Review: Self-Service Password Reset Managers p. 57

# Windows IT Pro

A PENTON PUBLICATION

JANUARY 2012 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

# vSphere 5 Delivers Scalability and Automation p. 21

LDAP Over SSL

Lock Down Active Directory Traffic p. 29

Deploy Exchange Server 2010 Personal Archives p. 34

Use PowerShell to Report on Scheduled Tasks p. 41

On-Premises vs. Hosted Email Archiving p. 45

10 Reasons Not to Brand SharePoint p. 48

Trends in Mobile Device Security p. 62

Interview: Jeffrey Snover, Lead Architect for Windows Server p. 24

# Netezza. Up and running in 24 hours, not 24 days.

Get set up in hours instead of days, and start counting returns in minutes instead of hours. All with IBM's Netezza data warehouse appliance for high-performance analytics. It gives you analytics reports at supersonic speeds. At a fraction of the cost of Oracle Exadata. Get real, actionable business results fast.

ibm.com/facts

COST comparison based on publicly available information as of 2/10/2011 for an Oracle Exadata X2-2 HP Full Rack and a full rack of Netezza TwinFin. The cost to acquire Netezza can be as low as 1/6 of Exadata if a client is acquiring new Oracle database licenses and as low as 1/2 if using existing Oracle database licenses. 

# C O N T E N T S

Windows IT Pro
A PENTON PUBLICATION
JANUARY 2012
VOLUME_18 NO_1

## COVER STORY



## FEATURES



## INTERACT



## IN EVERY ISSUE



## COLUMNS

JAMES | IT PRO PERSPECTIVES



THURROTT | NEED TO KNOW



MINASI | WINDOWS POWER TOOLS



OTEY | TOP 10



SPRINGSTON | WHAT WOULD MICROSOFT SUPPORT DO?



Access articles online at www.windowsitpro.com. Enter the article ID (located at the end of each article) in the Search box on the home page.

# CONTENTS

## PRODUCTS




Windows IT Pro

### EDITORIAL

**Editor in Chief**
Amy Eisenberg amy@windowsitpro.com

**Senior Technical Director**
Michael Otey motey@windowsitpro.com

**Technical Director**
Sean Deuby sean@windowsitpro.com

**Senior Technical Analyst**
Paul Thurrott paul@windowsitpro.com

**Industry News Analyst**
Jeff James jjames@windowsitpro.com

**Custom Group Editorial Director**
Dave Bernard dbernard@windowsitpro.com

**Exchange & Outlook**
Brian Winstead bwinstead@windowsitpro.com

**Systems Management, Networking, Hardware**
Jason Bovberg jbovberg@windowsitpro.com

**Security, Virtualization**
Jeff James jjames@windowsitpro.com

**SharePoint**
Caroline Marwitz cmarwitz@windowsitpro.com

**SQL Server, Developer Content**
Megan Keller mkeller@windowsitpro.com

**Managing Editor**
Lavon Peters lavon.peters@penton.com

**Editorial Assistant**
Blair Greenwood blair.greenwood@penton.com

### CONTRIBUTORS

**SharePoint and Office Community Editor**
Dan Holme danh@intelliem.com

**Senior Contributing Editors**
David Chernicoff david@windowsitpro.com
Mark Minasi mark@minasi.com
Paul Robichaux paul@robichaux.net
Mark Russinovich mark@sysinternals.com
John Savill john@savilltech.com

**Contributing Editors**
Alex K. Angelopoulos aka@mvps.org
Michael Dragone mike@mikerochip.com
Jeff Fellinge jeff@blackstatic.com
Brett Hill brett@iisanswers.com
Darren Mar-Elia dmarelia@windowsitpro.com
Tony Redmond 12knocksinna@gmail.com
Eric B. Rux ericbrux@whshelp.com
William Sheldon bsheldon@interknowlogy.com
Curt Spanburgh cspanburgh@scg.net
Orin Thomas orin@windowsitpro.com
Douglas Toombs help@toombs.us
Ethan Wilansky ewilansky@windowsitpro.com

### ART & PRODUCTION

**Production Director**
Linda Kirchgesler linda@windowsitpro.com

**Senior Graphic Designer**
Matt Wiebe matt.wiebe@penton.com

### ADVERTISING SALES

**Publisher**
Peg Miller pmiller@windowsitpro.com

**Director of IT Strategy and Partner Alliances**
Birdie J. Ghiglione birdie.ghiglione@penton.com
619-442-4064

**Online Sales Development Director**
Amanda Phillips amanda.phillips@penton.com

**Key Account Director**
Chrissy Ferraro christina.ferraro@penton.com
970-203-2883

**Account Executives**
Barbara Ritter barbara.ritter@penton.com
858-367-8058
Cass Schulz cassandra.schulz@penton.com
858-357-7649

**Client Project Managers**
Michelle Andrews 970-613-4964
Kim Eck 970-203-2953

**Ad Production Supervisor**
Glenda Vaught glenda.vaught@penton.com

### MARKETING & CIRCULATION

**Customer Service** service@windowsitpro.com

**IT Group Audience Development Director**
Marie Evans marie.evans@penton.com

**Marketing Director**
Sandy Lang sandy.lang@penton.com

### CORPORATE

PENTON

**Chief Executive Officer**
Sharon Rowlands sharon.rowlands@penton.com

**Chief Financial Officer/Executive Vice President**
Nicola Allais nicola.allais@penton.com

### TECHNOLOGY GROUP

**Senior Vice President, Technology Media Group**
Kim Paulsen kpaulsen@windowsitpro.com




**WRITING FOR *WINDOWS IT PRO***
Submit queries about topics of importance to Windows managers and systems administrators to articles@windowsitpro.com.

**PROGRAM CODE**


**LIST RENTALS**
Contact MeritDirect, 333 Westchester Avenue, White Plains, NY or www.meritdirect.com/penton.

**REPRINTS**
Wright's Media
penton@wrightsmedia.com 877-652-5295

# *Windows IT Pro* Congratulates
# Solarwinds

solarwinds


Category: **Network Management**
Product: **Network Performance Monitor**
Award: **Community Choice Gold**


Category: **Network Management**
Product: **Network Performance Monitor**
Award: **Editor's Best Gold**

Learn more about Solarwinds here:
http://www.solarwinds.com/?CMP=SYN-TAD-WITPRO-NPM_AWARD-NPM-SWHP-GOLD_Q4_2011
• 866.530.8100

# James

"Windows 7 is a superior OS in just about every way possible, with better security, reliability, and manageability, and a host of other real and substantial improvements over Windows XP."

# Long Live Windows XP

## Windows 7 is unarguably a better OS, but financial constraints and compatibility issues are keeping many IT departments on Windows XP

Microsoft's recent 10-year anniversary for Windows XP was a widely heralded event, and accolades freely flowed for Microsoft's longest-lived OS ever. Sadly, Microsoft's efforts at highlighting the commendable track record of XP were overshadowed by overly zealous efforts by various Microsoft spokespeople to drive customers off of XP and onto Windows 7. Like the scene from *Monty Python and the Holy Grail,* when the hapless peasant claims he isn't quite dead yet, Microsoft seems all too eager to boot XP into the grave in favor of Windows 7.

Granted, Windows 7 is a superior OS in just about every way possible, with better security, reliability, and manageability, and a host of other real and substantial improvements over XP. Recent data from StatCounter shows that, as of October 2011, Windows 7 had more market share than XP. But Microsoft clearly wishes people would move off of XP more quickly.

As Microsoft has continually added new features to Windows over the years, it seems like more IT departments than ever are starting to drag their feet when it comes to client OS upgrades. The evolution of Windows that resulted in XP created a product that, to many IT managers and CFOs, still passes the "good enough" test 10 years after launch. With businesses cutting costs, trimming fat, and laying off workers to stay afloat in harsh financial times, once thorny financial decisions become razor sharp: Do we upgrade to Windows 7 and lay off half the people in marketing, or do we soldier on with XP for a few more years?

More IT departments than ever are starting to drag their feet when it comes to client OS upgrades.

Virtualization, replacements for failing hardware, and Software as a Service (SaaS) solutions might be more attractive IT upgrades to cash-strapped management than an upgrade to a client OS that's still largely getting the job done. It's like having an old beater of a pickup truck that has 125,000 miles on the odometer, rust spots, and threadbare seat covers. It gets you where you need to go, so why spend more money on a new truck with hefty insurance premiums and expensive monthly payments? For better or worse, many IT departments are reluctantly clinging to XP for the same reasons.

A few informal queries I directed at IT pros on Twitter over the past few months resulted in some interesting replies about why they haven't dumped XP yet. Twitter user @Twirrim said that it's "hard to justify buying Win 7 licenses for the XP machines, esp when XP 'just works' and is less bloated," while @DecHL said, "Why still XP? Boss is too cheap to allow me to refresh the remaining boxes on XP. Got most machines running Win7." And @prudentdad echoes the sentiment of several other IT pros who are held back by XP or Internet Explorer 6 (IE6) compatibility requirements with key applications: "Still using Windows XP due to our document management software."

This lugubrious migration off of XP might be a sign of more ominous things to come for Microsoft's venerable product strategy of releasing a new Windows client every few years. Despite Microsoft's best efforts at cajoling, prodding, and soon shoving (via the April 8, 2014, end-of-life deadline) IT departments off of XP, what happens when the majority of computing is done with tablets and smartphones? iPhone, iPad, and Android OS users have become accustomed to mobile OS software updates and upgrades being quick, largely painless, and free.

Windows 8 tablets will more than likely put a huge dent in Apple's dominant tablet market share, but I predict that Microsoft never replicates the more than 90-percent market share it enjoys on the PC desktop in the tablet or smartphone market. My take is that those days are long gone, replaced by a multi-polar world where tablets, smartphones, and desktop PCs distribute workloads evenly.

A deeply troubling proposition for Microsoft is this one: If more than 50 percent of computing is done on devices where OS upgrades are free, what happens to the Windows business model when the OS moves to tablets and smartphones? Billions in lost revenue, that's what.

InstantDoc ID 141341

**JEFF JAMES** (jjames@windowsitpro.com) is industry news analyst for *Windows IT Pro*. He was previously editor in chief of Microsoft *TechNet* magazine, was an editorial director at the LEGO Company, and has more than 15 years of experience as a technology writer and journalist.

- Exchange Autodiscovery
- VSS Backup
- Windows 8
- Cloud and PowerShell



# LETTERS@WINDOWSITPRO.COM

## Exchange Autodiscovery

I really like *Windows IT Pro* and find most of the articles very informative and on track, which is why I was surprised to read John Savill's answer to the question "How can I quickly verify that my Exchange autodiscovery is working?" (InstantDoc ID 139558). The answer he provided, I think, not only doesn't accurately answer the question but also neglects to mention probably the best way to check full autodiscovery functionality on all devices. When it is set up this way, additional SSL certificates aren't necessary and it will work with multiple domain names.

For me, the correct (and Microsoft-recommended) answer is documented in the Microsoft article "A new feature is available that enables Outlook 2007 to use DNS Service Location (SRV) records to locate the Exchange Autodiscover service" (support.microsoft.com/kb/940881). This method doesn't require an additional SSL certificate. After completing the instructions in the article, you should go to the Microsoft Remote Connectivity Analyzer (www.testexchangeconnectivity.com) to test, verify, and help troubleshoot any Autodiscover problems.

—Matt McHugh

*Thanks for writing! The FAQ you reference is for checking only whether the autodiscovery component is responding and functioning rather than checking full connectivity. For checking full connectivity, see my followup FAQ, "How can I check on the health and accessibility of my organization's ActiveSync service?" (InstantDoc ID 141374).*

—John Savill

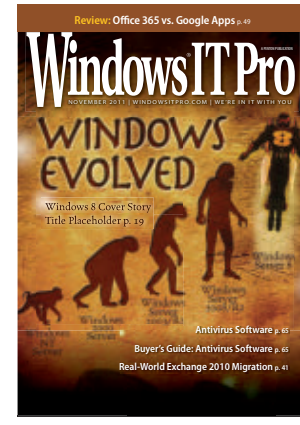


## Old-School Backups vs. VSS

In the first paragraph of his FAQ "What command can verify that the Hyper-V Volume Shadow Copy Service (VSS) writer is registered?" (InstantDoc ID 139638), Greg Shields describes old-school backup procedures. He writes, "Any changes that occur during the backup are then incorporated at the job's completion." VSS backups don't operate this way. When VSS signals quiescence, every block that changes is written elsewhere, leaving the old block intact. The old blocks are used by the backup procedure; the new blocks are used for the normal operation of the database. When the backup finishes, it doesn't incorporate any of the changes that occurred during its operation. This is desirable. With VSS, we obtain point-in-time backups—the point in time being the beginning of the backup procedure.

—Dimitrios Kalemis

*A bit of research suggests that you might be correct, although the "correctness" is ambiguous. VSS itself doesn't facilitate incorporating any changes to a backup at job completion, although a backup agent might (even though this is highly doubtable). I stand corrected and have removed the sentence.*

—Greg Shields

## Phone-ification of Windows 8

I realize the purpose of *Windows IT Pro* is, well, Windows—but "The Wonderful

*Windows IT Pro* welcomes feedback about the magazine. Send comments to letters@windowsitpro.com, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.

Phone-ification of Windows 8" (InstantDoc ID 140642) is too much. Paul's gushing love of Windows Phone is fine. He is, after all, paid to have an opinion, and at least he admits to being in a tiny minority that loves the phone enough to not only buy one but also write a book on the subject. But remember that Windows Phone isn't a 1.0 product. Microsoft has been in the phone business for years. Long before Apple got into the business, before Android was even conceived, Microsoft was putting Windows on phones that few people bought. That history of missing the mark says something. And let's not even start about tablet computing—a miserable failure until Apple showed people how to do it.

The new interface is interesting, to be sure. It seems actual users were consulted this time. But to a casual observer, it's not that different from the competition. It is nice that "you just go to the Photo hub" to look for pictures; a data-centric view rather than a programmatic view is something we should have had long ago. But it might be instructive to explore why the world is ignoring Windows Phone. Could it be the reportedly buggy OS? The lack of (at this point at least) multi-tasking? Poor performance? Or is it the lack of continuity? Each Windows smartphone incarnation has been heralded as "brand new," or "redesigned from the ground up." A few rounds of that makes a consumer wonder what was so bad with the previous version that the company was forced to throw it out and start over—to say nothing of replacing applications after buying the new version. And where are those applications? Those take time to develop, and with a history of OS replacement, it will be hard to get developers to commit to the new platform.

I suspect, however, that the biggest problem with sales is the "Me too!" syndrome. The new interface for Windows 8 looks a lot like what Apple has done with OS X Lion—a good move, but Apple released its OS two months ago while Windows 8 is a year or more out. Likewise, the blurred line between mobile and desktop computing, and the apps store—you get the point. It's not that Microsoft is heading in the wrong direction. It's just that it appears to be following the crowd. If Microsoft can catch up, it might gain some market share. But to actually lead in the mobile market, it has to completely change its business model for rapid development. It will need to develop useful new ideas that can be implemented now—ideas that consumers actually want. And it will need to pick a direction and stick to it. Otherwise, people will ignore the new hotness for something that works.

—Randy Grein

## Cloud, Cloud, Cloud . . . and PowerShell

I completely agree with the comments made by Stoney Heflin on the November *Windows IT Pro* Community Forum page (InstantDoc ID 140714). Cloud, cloud, cloud—every time I read about this great vision of the cloud, I wonder how many IT managers are ready to stake their careers on that vision. I, for one, am not. I realize that you have two different audiences to please: those who are pushing the cloud (the sellers) and those who have to decide to use it (the buyers). But please try to present the stories as unbiased as possible. Until cloud computing is as trustworthy as all the sellers want you to believe, every article dealing with the cloud should come with serious disclaimers as to its reliability.

Does anyone remember Novell NetWare? I do. I installed the first NetWare 3 file server at our company. In fact, we're still running a NetWare 3 server! It just runs, and runs, and runs. We restart it every few months. My only complaint with NetWare was the plethora of commands required to configure and manage it. Then along came Windows. Wow, what a change! Instead of a thousand commands to remember, everything was handled with a point-and-click GUI. Out with NetWare, in with Windows. Little did I realize that the time would come when Microsoft's engineers would forget how to make a GUI. Or perhaps its servers are too complex to manage with point-and-click. Now we have to use PowerShell to manage Windows. That's bad enough, but even worse is the fact that Microsoft markets PowerShell as the greatest innovation since sliced bread. I realize that complaining about PowerShell will be as effective as spitting in the ocean. I can only guess that Microsoft hopes that, someday, Windows will be as good as NetWare. Because—thanks to PowerShell—it's getting just as difficult to manage.

—Alf Flowers

I completely agree with the comments from reader Stoney Heflin. I have no plans to move anything to the cloud in the next five years and would appreciate more articles focused on servers, applications, and administrative procedures.

—Robert Matthews

InstantDoc ID 141372

**Instant Poll Results: Exchange Management Shell**

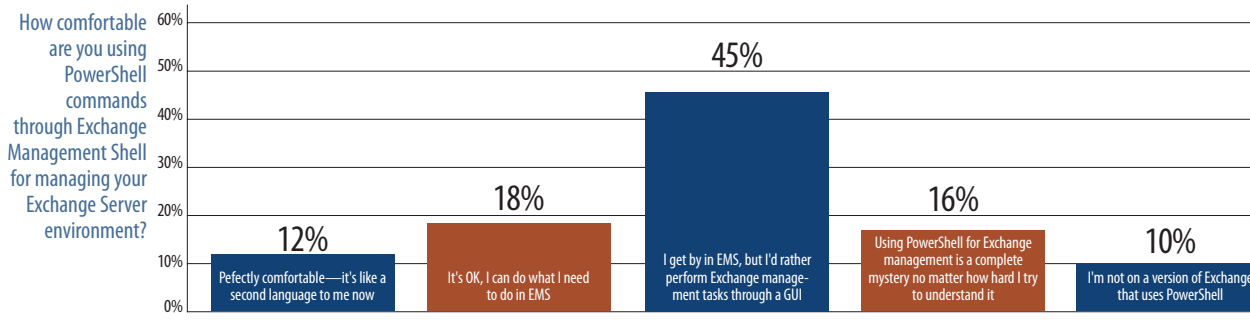


Source: *Windows IT Pro* Instant Poll, www.windowsitpro.com, November 2011

# Thurrott

"Windows Phone 8 could simply be a highly modified version of mainstream versions of Windows 8. Heck, Apple did it with iOS."

# The Rise of Windows Phone, Windows 8 Beta Predictions, and How Microsoft Needs to Manage Its Smartphone's Future

With a new year come changes: new product releases, technological shifts, and evolving expectations. And 2012 will be no different. This year, we'll see the release of Windows 8 (unless something goes horribly wrong), giving Microsoft a toehold in the crucial tablet market. This year, Windows Phone will establish itself as the clear number three in the handset market, or it will simply fade away. And this year, Microsoft will edge further from its roots in traditional software markets and move inexorably to the cloud and connected services. I can't wait to see what 2012 brings.

## Evidence of the Rise of Windows Phone...

It's going to be a while before unit sales of Windows Phone handsets are strong enough for Microsoft to begin calling them out in financial reports. As of the end of 2011, Windows Phone was racking up single-digital market share at best, well behind Google Android, Apple iOS, and Research in Motion (RIM) BlackBerry, and, embarrassingly, even behind the nebulous "other" category. But with a new generation of software and capabilities (Windows Phone 7.5), new handsets (including the reintroduction of Nokia to the US market), and missteps by competitors, Windows Phone, finally, might be gaining a foothold.

It starts with developers. Looking back on the first year of Windows Phone, I see perhaps one major success story: the platform's developer story. Microsoft made a lot of right decisions around the Windows Phone platform, basing the developer tools on existing languages, frameworks, and environments that are already familiar to programmers. This means C#, Visual Basic, and other familiar .NET programming languages; Microsoft Silverlight and XNA frameworks, which are familiar to anyone who's written apps targeting Silverlight, Windows Presentation Foundation (WPF), or other .NET frameworks on Windows; and Visual Studio (VS), hands-down the best development environment, period.

Microsoft has also done an exemplary job courting the right developers. As I write this, the total Windows Phone app count is at about 40,000 apps—not too shabby for a platform that's been around for all of 15 months. And of the top key apps on both Android and iOS, over 90 percent are available on Windows Phone. This includes heavyweights such as Facebook, Netflix, Evernote, Spotify, and YouTube, and of course games such as Angry Birds and Plants vs. Zombies.

Microsoft has also gone out of its way to provide reams and reams of Windows Phone developer documentation, which takes the form of (often free) eBooks, printed books, web-based documentation, and videos. An amazing array of ever-growing content is out there, aimed at developers of all levels.

This effort is paying off. A joint survey of mobile developers by IDC and Appcelerator claims that Windows Phone has "decisively" moved ahead of RIM's BlackBerry OS to become the clear number-three mobile OS behind Android and iOS. This is thanks to four major changes in 2011: platform improvements in Windows Phone 7.5 (the curiously named second major release of Windows Phone), the full backing of cell phone giant RIM, the ongoing "collapse" of BlackBerry, and HP pulling the rug out from underneath its Palm webOS platform.

"Windows Phone 7 separated from the pack to become the clear number-three mobile platform this quarter," the report reads. "The OS climbed 8 points to 38 percent of respondents saying they are 'very interested' in the platform, the highest ever for Microsoft." Granted, the companies have only been tracking this information for two years. But by comparison, the iPhone is targeted by 91 percent of mobile developers and Android handsets rack up 83 percent of the market.

BlackBerry, as expected, is dropping off the face of the earth. The platform fell 7 points to just 21 percent interest, and as the report notes, even Nokia's new Windows Phone–based Lumia handsets garner more interest among mobile developers than does BlackBerry.

## ...But Still a Lot of Work Ahead

Of course, with Windows Phone, it's hard to write anything positive without introducing some caveats. Indeed, many questions arise.

The first concerns Windows 8. With Windows 8 adopting and enhancing Windows Phone's Metro-style UI, it's hard not to imagine the two platforms coming closer together. In fact, when you consider the hardware that runs at the heart of Windows Phone handsets—1GHz to 1.5GHz ARM-based processors, hardware-accelerated graphics, 512MB to 1GB of RAM, 8GB to 32GB of storage—you see that these highly mobile devices are, in many ways, tiny PCs. So one's mind turns to the notion that Windows Phone 8

could simply be a highly modified version of mainstream versions of Windows 8, bringing the two platforms to a logical combination. Heck, Apple did it with iOS, which is a modified version of OS X.

I've heard from two reliable sources at Microsoft that that's exactly what the software giant is planning. There's no public confirmation on this rumor, so take it with the proverbial grain of salt. But I think it's happening. In fact, I think it's inevitable.

One way or the other, Microsoft will have to offer minor Windows Phone updates in the first half of the year, a release some are calling "Tango," which seems to be largely about making Windows Phone more accessible to the lower end of the market. (This part of the market is less expensive and larger from a volume standpoint than the traditional smartphone market.) After that, there will be a major release as well. With Windows 8 expected in the second half of 2012, a Windows Phone 8 release around the same time does make some sense.

There are compatibility issues to consider around such a transition, not to mention developer-related concerns, since the underlying runtime on Windows 8, called WinRT, isn't the same as that in today's Windows Phone OS. So developers would need to make yet another transition, which might not be as horrible as it sounds, given the language/framework/environment similarities between the two platforms. But it's still a transition.

Regardless of what Microsoft does with the "guts" of Windows Phone, I see a pressing need for the company to conceptually advance the Windows Phone platform in some key areas in 2012. That includes removing its reliance on the terrible Zune PC software, which hasn't been updated in a significant fashion since well before Windows Phone. I recommend that Microsoft get rid of this weird dependency and build Windows Phone connectivity directly in Windows 8 for those few people who need such things (as well as for those who need to seamlessly download phone-based photos to the PC, another area where the Zune PC software comes up limp).

And I do mean those few people. Smartphones have evolved from PDA-like PC companions and are now full-fledged portable computing devices in their own right. We need to evolve with them, moving the hub for our personal- and work-related data from the PC to the cloud. In such a system, the PC becomes an endpoint for that data, just like a phone or tablet.

In Windows Phone today, some areas that previously needed PC connectivity for accounts management (email, contacts, calendar) have already moved to the cloud, but other areas—media management, primarily, and photo downloading—still require PC connectivity. Removing these ties will put Windows Phone on par with competing platforms such as iOS—which uses the iCloud services for media management and photo downloading, among other things—and make the platform more universally relevant. Supporting PC users and Mac users would be basically identical. And resetting a device would be less painful because everything on that device is stored in the cloud.

I've been told by sources that Microsoft does intend to better integrate Windows 8 with the next versions of Windows Phone and the Xbox console, though what form that integration will take is still unclear. I've heard terms like "embedded Silverlight" for the next Xbox, code-named TEN, and "Apple-like integration" between the various pieces. We'll see.

## Windows 8 Beta

Microsoft is expected to deliver the first (and only) beta version of Windows 8 at or around the 2012 Consumer Electronics Show (CES) January 10–13. This release is notable for many reasons, not the least of which is that it's only one of an expected three major prerelease milestones for the product. Whether Microsoft is able to deliver this crucial beta release at CES will say a lot about the coming schedule for the final version of the product.

Windows 8, as you know, is a big bet for Microsoft and a potential disaster in the making if consumers and business customers don't embrace its new "touch-first" usage model. Yes, it should be close to fully compatible with Windows 7–era applications and hardware. But it's the forward-leaning Metro-style UI, the Start screen, and the underlying Windows Runtime that will make or break it.

Last year, I was told to expect a feature-complete version of Windows 8 at what became the BUILD Conference. But the prerelease Developer Preview wasn't feature-complete and provides only a peek at the full Windows 8 user experience. The key here, as I discussed in "The Windows 8 Paradox, A Mobile Market Reshuffle, and RIM's Nosedive Into Obscurity" (www.windowsitpro.com/article/windows8/windows-8-paradox-mobile-market-rims-nosedive-140998), is that it's jarring moving back and forth between new, Metro-style apps and classic, legacy desktop applications. But since all we have today are the legacy applications, we're not getting the full Windows 8 experience.

What's going to change isn't just the platform improvements that Microsoft brings to the beta, but also the opening of the Windows Store, where Microsoft and third-party developers will sell and give away new, Metro-style apps. My understanding is that the Windows Store could open for business as early as the launch of the Windows 8 beta in January, meaning that we will be able to more reasonably assess how well this new OS will perform in the real world.

Also up in the air is the status of the ARM-based versions of Windows 8. Microsoft is characteristically coy about how it will implement these versions, but I expect that ARM-based Windows 8 versions will appear exclusively (or nearly so) on iPad-like tablet devices and not provide any access to the legacy Windows desktop or its applications. That would leave more traditional x86/x64 versions of the OS to a wider range of "true" PCs, which would include desktops, laptops, tablet PCs, and slates. I hope Microsoft will clear this up early in 2012 as well.

Whatever happens, 2012 is going to be a huge year for those who are interested in Windows 8. Microsoft's next OS will ship, one way or another, and I expect we'll be dealing with the repercussions of that for some time to come. And that will be particularly true if Windows 8 doesn't catch on as well as Microsoft expects.

InstantDoc ID 141272

**PAUL THURROTT** (paul@windowsitpro.com) is the senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows (winsupersite.com), a weekly editorial for *Windows IT Pro UPDATE* (www.windowsitpro.com/email), and a daily Windows news and information newsletter called *WinInfo Daily UPDATE* (www.wininformant.com).

# DB2 on POWER: 3x faster. Check. As low as 1/3 the price. Mate.

Which database has the right moves? DB2® on Power Systems™ performs three times faster per core than Oracle Database on SPARC—based on both TPC-C and SAP® SD benchmarks.* Yet the price of DB2 is as low as 1/3 the price of Oracle Database.** Maybe that's why in 2010 over 1,000 Oracle Database clients chose DB2 instead. Game over.

ibm.com/facts

*PERFORMANCE: www.tpc.org as of 3/28/11 [IBM Power 780 (3 x 64 C)(24 Ch/192 C/768 Th); 10,366,254 tpmC; $1.38/tpmC; avail. 10/13/10 v. Oracle SPARC SuperCluster w/T3–4 Servers (27 x 64 C)(108 Ch/1728 C/13824 Th); 30,249,688 tpmC; $1.01/tpmC; avail. 6/1/11]. TPC-C is a trademark of Transaction Performance Processing Council. 2-tier SAP SD standard application benchmark results as of 3/28/11 [IBM Power 795 (32 P/256 C/1024 Th); 126,063 users, SAP ERP 6.0 EhP4/AIX 7.1 + DB2 9.7; cert. 2010046 v. Oracle SPARC Enterprise Server M9000 (64 P/256 C/512 Th); 39,100 users, SAP ERP 6.0/Solaris 10, Oracle 10g; cert. 2008042] www.sap.com/benchmark. SAP and all SAP logos are trademarks or registered trademarks of SAP AG in Germany and several other countries. **PRICE: based on publicly avail. U.S. info on 2/10/2011 for IBM DB2 Advanced Enterprise Edition + Oracle software w/comparable capabilities. No SAP SD benchmark results are used for any price/performance metrics. IBM: 100 Processor Value Units. Oracle: assumes 1.0 processor multiplier. Both incl. Y1 maint./support. 

# Minasi

"Trying to pick out a subset of users according to a criterion? Get-ADUser is the place to start."

# Use Get-ADUser to Determine Who Has Never Logged On

## Learn how to craft a more interesting AD query

Through the past few columns, I've been demonstrating Windows Server 2008 R2's useful Active Directory (AD)-related PowerShell cmdlet Get-ADUser. So far, though, my sample Get-ADUser queries have been pretty trivial. Now it's time to build a more interesting query by asking AD, *Which accounts have never logged on?* In the process, you'll improve your PowerShell skills.

Whenever you're trying to pick out a subset of your users according to some criterion, Get-ADUser is the place to start. But to make it work, you have to build a filter, and that filter will almost certainly refer to one or more of the 100-plus attributes that every AD user's account contains. The tough part is usually figuring out which of those attributes can assist you, so let's explore what the list of attributes offers. You can see the attributes by typing

```
get-aduser -filter *-properties *|get-member
```

Searching through the attributes that *get-member* provides for ones that have *logon* in their names, you'd find the most promising ones to be LastLogon, LastLogonDate, LastLogonTimestamp, and LogonCount. MSDN's AD schema documentation shows that LastLogon and LastLogonDate won't be useful, as they're not replicated amongst DCs. LastLogonDate isn't even in the schema, so we'll focus on LastLogonTimestamp. (In truth, I'm cheating. Most AD geeks have known about LastLogonTimeStamp since 2003, but that's how I'd have researched it, had I just come to the AD party.) Having found this candidate attribute, then, the next step would be to see a few of its typical values with this command:

```
get-aduser -f * -pr lastlogontimestamp | ft samaccountname,
  lastlogontimestamp -auto
```

Here, I'm telling Get-ADUser to retrieve all users but to pipe the output to *ft* (short for *format-table*), which I instruct to show only the *samaccountname* and *lastlogontimestamp* values, while making the table as compact as possible with the *-auto* parameter. Its output, when run on a little test AD environment, looks like

```
samaccountname lastlogontimestamp
-------------- ------------------
Administrator 129699195295312500
Guest
krbtgt
julesm 129699193327187500
```

This reveals something interesting about *lastlogontimestamp:* It shows up either as a huge number or as nothing. Guessing that "nothing" equals "has never logged on," I tested that by creating a new user account without logging it on, ran the Get-ADUser query again, and it too had no *lastlogontimestamp*. My query, then, need only look for accounts with empty *lastlogontimestamp* values.

Unfortunately, PowerShell doesn't make that easy. You can't just use a filter like {lastlogontimestamp -eq ""} or {lastlogontimestamp -eq $null}; instead, you have to use *-like* (which we've met already) and *-not* (which we haven't) to write the filter this way:

```
get-aduser -f {-not ( lastlogontimestamp -like "*")}
```

The *lastlogontimestamp -like "*"* parameter matches any record where there's anything at all in the *lastlogontimestamp* attribute, essentially finding only the "non-empty-LastLogonTimestamp" records. The *-not* operator negates whatever you feed it—in this case, converting the "find all the non-empty attribute records" query to a "find all the empty attribute records" query.

Run that query, and you'll see all the never-logged-on folks, but you'll see that the query also returns built-in accounts such as krbtgt and Guest—things that you'll never want to mess with and don't want in your query results—so you'll need to refine the query filter. A look at krbtgt and Guest reveal that they're disabled. Thus, you need to change the query from "choose all users with an empty lastlogontimestamp" to "choose those users with that empty attribute and that aren't disabled." That query looks like

```
get-aduser -f {-not ( lastlogontimestamp -like "*") -and
  (enabled -eq $true)}
```

The only difference is that I've added *-and*, as well as another filter, *enabled -eq $true*. That filter means "isn't disabled" because AD's PowerShell tools give users an attribute named *enabled* rather than *disabled*, and because PowerShell represents the logical value *true* with the phrase *$true*.

In this example, you've seen how to explore an AD object's structure and common values. You then built more complex queries by detecting empty attributes, flipping query filters on their head with *-not* logical values, and using *-and* to join multiple filters.

InstantDoc ID 141189

**MARK MINASI** (www.minasi.com/gethelp) is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books.

**Otey**

"Multi-touch, along with the Metro UI and a new generation of touch-enabled devices, could make Windows 8 an iPad killer."

# New Features in Windows 8

## The Metro UI, multi-touch support, and new power management make Windows 8 ready for tablets and the desktop

At the inaugural BUILD 2011 conference in Anaheim this past September, Microsoft unveiled the first official look at its upcoming Windows 8 desktop OS. Recent critics have rallied against Windows for being long in the tooth, saying the PC is a dying form factor. The Windows 8 features that Microsoft showed at BUILD make these statements seem way off base. Windows 8 is so feature-packed that I can't come close to listing them all in a single column. However, here's what I think are the 10 coolest features in Windows 8.

1 **Metro UI—**Patterned after the Windows Phone interface, the tile-based Metro interface is the biggest UI change that Windows has ever seen. Groups of tiles populate the Start screen. Selecting a tile runs an app. An app toolbar slides up from the bottom of the screen and another toolbar containing Charms slides in from the right side of the screen. Charms let you change app-specific settings.

2 **Multi-touch support—**The Metro interface has a tile-based design, but the thing that really makes it exciting is the fact that it's multi-touch enabled. Multi-touch support lets the Windows 8 Metro UI run well on devices such as tablets and slates. Multi-touch, along with the Metro UI and a new generation of touch-enabled devices, could make Windows 8 an iPad killer.

3 **Aero interface—**For backward compatibility, Windows 8 includes the soon-to-be classic Aero interface, which was introduced with Windows Vista. In these early builds, the Windows 8 Aero interface is essentially identical to Windows 7 with a few minor changes, such as a new Task Manager. The Aero interface runs side-by-side with the Metro interface. Switching between them is instantaneous. The x86 and x64 versions of Windows 8 will be fully backward compatible with Windows 7 applications.

4 **Connected Standby power state—**Windows 8 changes the way Windows power management works. With Connected Standby, programs that aren't being actively used are put into suspended state and no code is executed. This system enables Windows 8 to run more efficiently on low-powered hardware and reduces the power footprint.

5 **Support for Hyper-V—**Like Windows Server 8 and the earlier Windows Server 2008, Windows 8 will include support for Hyper-V virtualization out of the box—there's no need for other virtualization programs such as Microsoft Virtual PC. Unlike Server 2008 R2 Hyper-V, the new Hyper-V 3.0 in Windows 8 is power-management friendly and will have the ability to suspend virtual machines (VMs) when the system goes to sleep or hibernates.

6 **Reset to default—**A new reset feature lets users reset their Windows 8 systems back to their default factory settings with one simple button click. Unlike the existing System Restore feature, the new reset option can remove all installed programs and user settings and it doesn't require that you perform any type of backup beforehand.

7 **Built-in antivirus software—**Another overdue feature is antivirus protection built in to the Windows 8 OS. Windows 8 will include the antivirus features from Microsoft Security Essentials. Additionally, if a compromised USB device is present at boot time, the system will refuse to start. These features might make antivirus vendors unhappy, but they'll go a long way toward making all Windows 8 systems more secure.

8 **Support for ISO and VHD files—**Because I work with both ISO and VHD file types frequently, this is one feature I'm particularly happy to see. Windows 8 can directly mount both ISO files and VHD files, letting you open and work with them directly by using Windows Explorer and other applications.

9 **Windows To Go—**Following in the footsteps of BitLocker To Go, Windows To Go will let you run a copy of Windows 8 that's booted from an external USB drive. Windows To Go is designed for use with flash drives having more than 32GB of storage. When you boot from a Windows To Go USB drive, the system's other physical drives are hidden.

10 **ARM support—**Another new feature in Windows 8 that's particularly important to running Windows 8 on tablets and other devices is support for ARM processors in addition to x86 and x64 processors. Support for ARM processors will open up an entirely new class of devices for Windows 8. However, existing x86 and x64 applications will need to be recompiled to run on the ARM platform.

InstantDoc ID 141257

**MICHAEL OTEY** (motey@windowsitpro.com) is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).

# Springston

WHAT WOULD MICROSOFT SUPPORT DO?

"The most difficult aspect of identifying an authentication bottleneck is that there's no event logged on any computer."

# Resolve Performance Problems Associated with Authentication Scaling

## Don't let legacy NTLM authentication get you down

Let's talk about legacy authentication on your network—specifically, scaling issues related to legacy authentication that cause you problems. Such problems can lead to outages once you reach a certain activity threshold. The concern boils down to the different performance characteristics of NTLM and Kerberos. The reason for preferring Kerberos authentication in Windows 2000 Active Directory (AD) and later isn't just the inherent enhanced security; it's also a better-*performing* authentication method. Each NTLM-based authentication is unique—even if it's a repeated authentication to the same resource by the same identity. Kerberos, on the other hand, provides a reusable access-granting service ticket for the resource to that identity, and that reuse requires no interaction with an authenticating server or domain controller (DC). NTLM is a more expensive authentication protocol, as well as a less secure one.

It's important to know when you've reached a point of authentication failure due to resource bottlenecks and an excessive volume of NTLM authentication. That's what this article is about: understanding that this problem is evident in your environment, and knowing how to fix it.

### The Details

A resource bottleneck can occur when a Windows computer needs to perform NTLM authentication for some user. (Figure 1 shows the NTLM authentication flow across domains.) For those familiar with the Windows architecture, you'll recall that the Local Security Authority Subsystem (lsass.exe) process is responsible for handling authentication requests. This is true for all versions and roles of Windows. There are threads within lsass.exe, and you might think of them as the workers that do the job of executing the code. For NTLM authentication, there's a maximum number of thread workers that can be running at any time to handle the job. The out-of-the-box defaults for that are to allow a single thread if the computer is a domain member and two threads if it's a DC. That NTLM thread is used on a domain computer to send a request to a DC, and a similar thread on the DC is used to craft the reply. So, in a typical transaction, there are at least two computers that can see this bottleneck. During that single authentication transaction of domain member to DC in that domain, the client thread is in a wait state until the DC replies.

If the user is requesting authentication from a trusted domain, you now have an additional DC contact to finish that authentication transaction. That wait state I mentioned before would now have the original client *and* that computer's DC tied up while waiting for the trusted DC's reply.

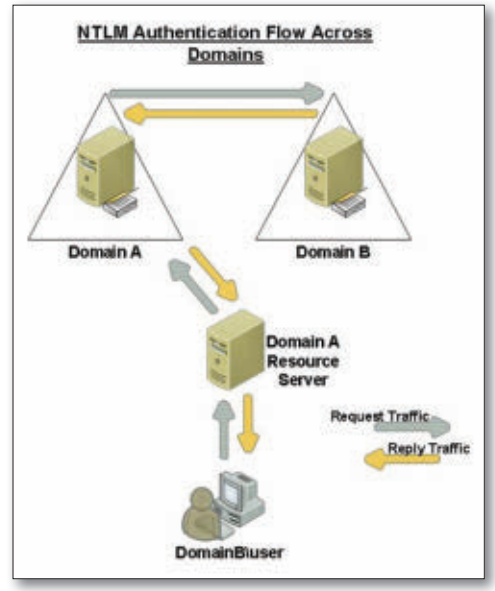


Figure 1: NTLM authentication flow

Of course, a thread executing an NTLM transaction is faster than the blink of an eye. Speed isn't a concern until you have a large number of *simultaneous* NTLM authentication requests, or if many of those transactions are across trusted boundaries to DCs in other domains. Add in a busy server—generating many NTLM authentication requests for its users—and you have a problem emerging.

The name of the limit on the NTLM authentications threads is MaxConcurrentApi. MaxConcurrent Api (of data type REG_DWORD) can be configured in the registry, under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters, and needs only the Netlogon service to restart to take effect.

MaxConcurrentApi is the Windows code that determines the creation of additional threads to handle new NTLM authentication requests. Without a thread to handle an authentication request, the requesting clients (which can be remote computers) might time out, become unresponsive, or return Access Denied errors to the user. That ambiguity is why it can be very difficult to figure out the root cause.

For all versions of Windows, the out-of-the-box default setting for MaxConcurrentApi is only 1 for a member server and 2 if the computer is a DC. In Windows Server 2003 and Windows Server 2008, you can change the MaxConcurrentApi setting as high as 10.

If you have Server 2008 R2, the maximum is 150, though the defaults are the same. If you have an original installation of Server 2008 (not R2), you can install a hotfix (described in the Microsoft article "You are intermittently prompted for credentials or experience time-outs when you connect to Authenticated Services" at support.microsoft.com/kb/975363), which will let you increase the maximum to 150 as well. That explains the mechanics of the bottleneck. Now let's talk about identification.

## Finding and Fixing the Problem

The most difficult aspect of identifying an authentication bottleneck is that there's no event logged on any computer. Instead, the errors all happen within the application that requested authentication. Depending on the application's error handling, there might not be enough details to pinpoint NTLM bottlenecks.

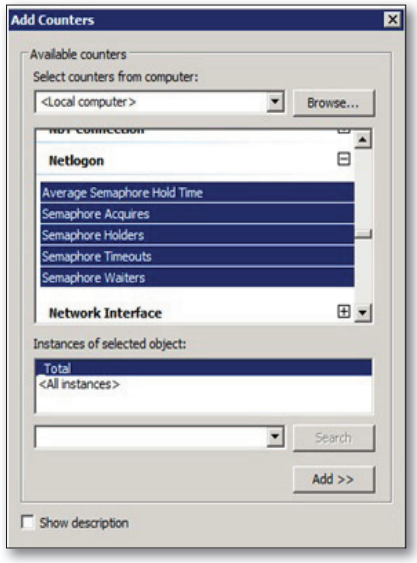


Figure 2: What to watch for in the performance log

Since you don't have an event and might not have a useful error, you need to look for other symptoms. The thing to keep in mind is that it can happen to any application using NTLM. Common culprits are old line-of-business (LOB) applications, which use NTLM because that was the lowest common denominator at the time.

The best way to tell whether you're reaching NTLM authentication bottlenecks is to determine if those failures are the result of volume. If the failures tend to occur during high-usage times (e.g., Monday morning, when users are arriving and beginning their work day), that's an indicator but not necessarily conclusive.

Use the Performance Monitor Netlogon performance object to monitor the server in question during a time when that server is under load. Note that you should do this on the resource server that users are having trouble accessing, as well as on the DCs; you don't want to miss a potential bottleneck. In the performance log (.blg) file, pay attention to the following (as Figure 2 shows):

- Semaphore Holders equal to the current value of the MaxConcurrentApi registry value setting
- Semaphore Timeouts with any number greater than 0
- Semaphore Waiters with any number greater than 0

If you have any timeouts or waiters, you have an NTLM authentication bottleneck.

Recall that I said trusted DCs might be involved and how that could increase those delays and timeouts. You can identify whether trusted domains are a factor by viewing this same performance data in the Report view, as Figure 3 shows. Each domain will appear with detailed numbers.

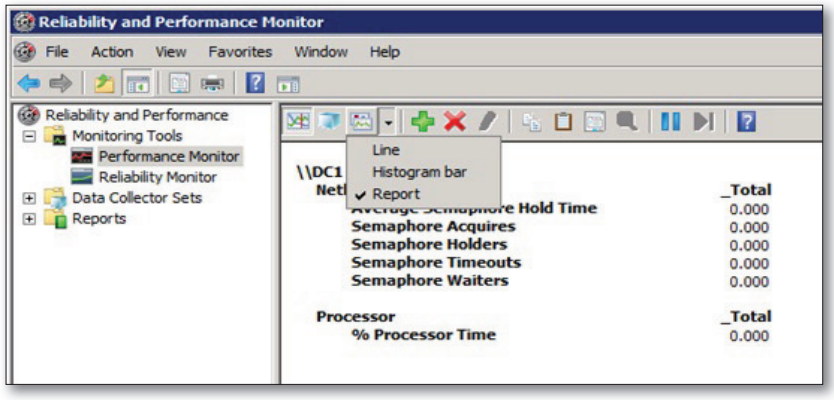


Figure 3: Performance data in the Report view

Identifying that you have the bottleneck is only the first step. Next, you need to address performance issues that are preventing your users from accessing services they need to get their jobs done. The easiest workaround is to increase the MaxConcurrentApi setting on all involved servers to a number that can handle the load. Because the maximum number is 10, it's best to raise it to 10 if you have Windows 2003 or Server 2008, or to a greater number if you have Server 2008 R2 (or the hotfix installed). Then, restart the Netlogon service on those servers.

When increasing the MaxConcurrentApi setting doesn't resolve the outage, you have to dig a little deeper to find out which computers and user accounts are sending the authentication requests. The Netlogon service debug log has those answers. (See the Microsoft article "Enabling debug logging for the Net Logon service" at support.microsoft.com/kb/109626 for more information.) This log isn't verbosely enabled by default, but it's easy to start, it won't fill up your drive, and it's time-indexed for reference.

Things to look for, both in the Netlogon service debug log and elsewhere—in the order of most common to least common—are the following:

- *NlpUserValidateHigher: Can't allocate Client API slot*—This text entry in the Netlogon log indicates that the computer has NTLM authentication requests waiting but is already at the maximum number of threads. The entries preceding this one will tell you the username and computer the request is coming from.
- *NlAllocateClientApi timed out*—This text entry in the Netlogon log indicates that one of the clients that was waiting to authenticate gave up after waiting 45 seconds. This entry's appearance means that a user somewhere received a credentials prompt, an error code, or an indefinite wait.
- *(null)*\—Null entries in the Netlogon log indicate that a legacy client on your network is submitting NTLM authentication requests for a domain user but omitting the domain of the user, so instead of domain\user you see (null)\user. In Windows 2003, this can result in extra use of those

authentication resources, therefore exacerbating a potential bottleneck into a real one. To resolve that concern, disable the ping behavior by using the Neverping setting, as the Microsoft article "The Lsass.exe process may stop responding if you have many external trusts on an Active Directory domain controller" (support.microsoft.com/kb/923241) describes. Note that this isn't a concern for Server 2008 and later.

- *Repeat offenders*—Frequent, repeated authentication attempts (i.e., entries that start with SamLogon) from the same user and computer appearing in the Netlogon log might indicate an application that is malicious or inefficient.
- *Kerberos PAC validation*—Oddly enough, this Kerberos security feature is implemented in Netlogon and uses those same threads that are a bottleneck for NTLM authentication. This behavior has an event that will appear in the System event log—event 7 with the source field of Kerberos. If you're seeing a high volume of these events and also seeing intermittent authentication outages for your users, try disabling this additional security feature temporarily until you can add more servers to handle the load. Disabling this feature permanently isn't recommended, and it's a moot point if it's an Exchange Server or IIS app pool service, since it cannot be disabled for them. Otherwise, the Microsoft article "You experience a delay in the user-authentication process when you run a high-volume server program on a domain member in Windows 2000 or Windows Server 2003 (support.microsoft.com/kb/906736) describes how to do it. If you confirm that you're seeing NTLM bottlenecks, the best solution is to use Kerberos instead. Older applications are less likely to support Kerberos, so that might not be an option. That can lead to some tough conversations, weighing the costs of budgeting for new software versus the need for security and scalability. Ultimately, outages and poor performance will help security and scalability win that debate every time.

## Trends

Two major trends are bringing this topic to the attention of IT folks everywhere: the consumerization of IT and the excellent performance of new hardware and software. Simply put, people want to use unmanaged or legacy clients to connect to really fast services over the cloud. It's your job to make sure those things "just work" for them—so take a good long look at your network environment and don't let authentication get you down.

InstantDoc ID 141269

**TIM SPRINGSTON** (tim.springston@microsoft.com) is a senior support escalation engineer in the Commercial Technical Support team at Microsoft, where he is the lead for security and authentication. Check out his Active Directory blog at blogs.technet.com/ad.

Custom PowerShell Function

# READER TO READER

## Emulating Cmd.exe's Set Command in PowerShell

I'm a long-time user of the command shell (Cmd.exe) in Windows. Old habits die hard, so I continually find myself typing the Set command in the Windows PowerShell console when I want to list, set, or clear environment variables. In PowerShell, Set is an alias for the Set-Variable cmdlet, but it doesn't work with environment variables. Instead, you have to use the Get-ChildItem, Get-Item, or Remove-Item cmdlet with the ENV: drive.

Bill Stewart

Rather than trying to break the habit, I decided to accommodate it by writing my own PowerShell Set function. I wanted my function to behave like Cmd.exe's Set command, as shown in Table 1. To get this behavior, my PowerShell Set function uses the MyInvocation object's InvocationName and Line properties to capture information about the function's command line. The InvocationName property contains the name of the command. The Line property returns the line used to invoke the command. Listing 1 contains the custom Set function.

You can't use the Set function as part of a PowerShell expression, such as

```
(Set processor_level).GetType()
```

The reason for this limitation is that the Set function needs to use a slightly unorthodox command-line parsing technique because the Line property returns the entire line, not just the expression. As a result, in this case, the *Environment variable not found* error is thrown.

However, the Set function has two advantages over Cmd.exe's Set command. First, it outputs DictionaryEntry objects, just like when you use the command

```
Get-ChildItem ENV:
```

Second, the Set function uses wildcard matching. For example, the command

```
Set P
```

matches only a variable named P, whereas in Cmd.exe this command will list all variables that start with P. To do this with the Set function, you'd use

```
Set P*
```

instead. This is equivalent to the PowerShell command

```
Get-ChildItem ENV:p* | Sort-Object Name
```

You can download the Set function's code (Set.ps1) by going to www.windowsitpro.com, entering 141086 in the Search box, and clicking the 141086.zip hotlink. To make the Set function a permanent part of your PowerShell command-line experience, put the code in your PowerShell profile.

—Bill Stewart, IT infrastructure group, Emcore

InstantDoc ID 141086

**Table 1: The Set Function's Behavior**

| Command | Result | Example |
|---|---|---|
| Set *name=value* | Sets the environment variable *name* to *value*. Quotes aren't required if *name* or *value* contain spaces. | Set NAME=Joseph Bogus |
| Set *name=* | Removes the environment variable *name* | Set NAME= |
| Set *string* | Displays environment variables starting with *string* | Set PROC |

**Listing 1: The Set Function**

```
# If an alias exists, remove it.
If (Test-Path ALIAS:set) { Remove-Item ALIAS:set }

Function Set {
  If (-Not $ARGS) {
    Get-ChildItem ENV: | Sort-Object Name
    Return
  }
  $myLine = $MYINVOCATION.Line
  $myName = $MYINVOCATION.InvocationName
  $myArgs = $myLine.Substring($myLine.IndexOf($myName) + $myName.Length + 1)
  $equalPos = $myArgs.IndexOf("=")
  # If the "=" character isn't found, output the variables.
  If ($equalPos -eq -1) {
    $result = Get-ChildItem ENV: | Where-Object { $_.Name -like "$myArgs" } |
      Sort-Object Name
    If ($result) { $result } Else { Throw "Environment variable not found" }
  }
  # If the "=" character is found before the end of the string, set the variable.
  ElseIf ($equalPos -lt $myArgs.Length - 1) {
    $varName = $myArgs.Substring(0, $equalPos)
    $varData = $myArgs.Substring($equalPos + 1)
    Set-Item ENV:$varName $varData
  }
  # If the "=" character is found at the end of the string, remove the variable.
  Else {
    $varName = $myArgs.Substring(0, $equalPos)
    If (Test-Path ENV:$varName) { Remove-Item ENV:$varName }
  }
}
```

Tell the IT community about the free tools you use, your solutions to problems, or the discoveries you've made. Email your contributions to r2r@windowsitpro.com.

*If we print your submission, you'll get $100.*

Submissions and listings are available online at www.windowsitpro.com. Enter the InstantDoc ID in the Search box.

- Outlook
- Hyper-V
- Office 365
- VMware
- vSphere 5

# ANSWERS TO YOUR QUESTIONS

## Q: What is Fair Share CPU Scheduling?

**A:** In OSs earlier than Windows Server 2008 R2, any user could make use of the full resources of a Remote Desktop Session Host. When resources were available, this relaxed distribution of resources didn't create a problem; however, when resources became scarce, the actions of one user could consume resources that were needed by others.

The Fair Share CPU Scheduling feature in Server 2008 R2 reduces the effect of users who make heavy use of resources by distributing processor time equally across the number of active sessions. With this feature in Server 2008 R2, the actions of one user can't affect the performance of the system for others, even under high load.

—Greg Shields

InstantDoc ID 141247

## Q: How do I configure Outlook and Exchange Server for use with Office 365 and Remote Desktop Services/Citrix?

**A:** Having the Exchange server in the cloud with Office 365 adds some interesting dynamics when your clients are using a virtual desktop environment, be it Virtual Desktop Infrastructure (VDI) or session based. Numerous guidelines and best practices exist around Outlook configuration and using its cached and online mode depending on the client's connectivity to the Exchange server. Typically, the guidance recommends using Outlook's Cached Exchange Mode when connectivity to the Exchange server isn't within the same local network, to ensure a great end-user experience.

However, it's the opposite if you have a virtual desktop environment and use Office 365. With a virtual desktop, it's best to keep the data small and use Outlook in its online mode, with no local cache. The good news is, Cached Exchange Mode is now supported in a remote desktop environment. However, the online mode is still the preferred mode.

The Microsoft document "Cached Exchange Mode in a Remote Desktop Session Host environment: planning considerations" (www.microsoft.com/download/en/details.aspx?id=15238) talks through the options. It will help you make the right decision for your organization.

—John Savill

InstantDoc ID 140934

## Q: I want to expand an Intel NIC driver; where does the extraction go when I add the /e switch?

**A:** I had this exact problem as I tried to extract an Intel NIC installation package so I could manually force a driver installation. I tried the /e switch, and it gave no indication of where the expansion had written to.

The secret is to also add the /f<path> switch, which tells the process where to write the extraction. Below I've added the /e and the /f<path> switches:

```
d:\temp\PROWinx64.exe /e /fd:\temp\
  intel
```

Note that there's no space between the /f and the path. For storage drivers, the /e switch doesn't work, so you need to use

```
-a -n -s -p <path>
```

after the storage driver file.

—John Savill

InstantDoc ID 141092

## Q: How does VMware vSphere 5's FDM select a master?

**A:** The master/slave architecture that vSphere 5's Fault Domain Manager (FDM) uses for monitoring vSphere HA clusters uses an election process to determine which host is to be the master. This election process occurs any time the existing master fails, is shut down, or is placed into maintenance mode. It also occurs when vSphere HA is enabled or when a management network partition occurs. The election process takes about 10 to 15 seconds.

The election process is defined by an algorithm with two published rules. For the first, the host with access to the greatest number of data stores wins. In the case of a tie, the second rule kicks in: The host with the lexically highest Managed Object ID (MOID) is chosen. Care must be taken when attempting to rig this election because lexically here means, for example, that host-99 is in fact higher than host-100.

—Greg Shields

InstantDoc ID 141016

Jan De Clercq | jan.declercq@hp.com
William Lefkovics | william@mojavemediagroup.com
John Savill | jsavill@windowsitpro.com
Greg Shields | virtualgreg@concentratedtech.com

## Q: What do I need to watch out for in managing the RID pool used in an AD domain? Or is this all done auto-magically?

**A:** In a Windows Active Directory (AD) domain, the process of generating unique Relative IDs (RIDs) is a single-master operation that's assigned to one specific domain controller (DC). This DC is then referred to as the RID master of the domain. The RID master can be hosted on either a DC or a Global Catalog (GC).

The RID master gives a pool of RIDs to each of the other DCs in the domain and keeps track of the sets of allocated RIDs for each DC. The domain-level RID pool controlled by the RID master can hold approximately 1 billion RIDs.

RIDs are never reused because the RID can't be reclaimed after a security principal is deleted. Reusing a RID could lead to unauthorized access to resources if the resources' access control settings referred to previously issued security IDs (SIDs) and RIDs.

The RID master gives every DC a pool of 500 RIDs at a time. When a new domain account or group is created, the DC assigns the new account a SID and a RID that's taken from its local allocated RID pool. When a DC's RID pool begins to run low, it automatically asks the RID master for another block of RIDs.

Problems occur when a DC has used all RIDs in its local RID pool and can't obtain a new pool from the RID master. For example, this could occur because of network problems.

The DC isn't able to create new security principals until a new local RID pool is obtained. In this case, event 16645 and possibly event 16651 will be logged in the Directory Services event log of the DCs that can't acquire new RID pools.

To reduce the chance of running out of RIDs, you can increase the number of RIDs that are allocated by the RID master to each DC's RID pool. To do this, adjust the RID Block Size value (REG_DWORD) on the RID master DC. The RID Block Size value is located in the following registry subkey: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\RID Values.

If you change the RID Block Size value, you should configure the new value not only on the RID master DC but also on all other DCs in your domain. That way, if the RID master needs to be transferred to another DC, the RID Block Size will be consistent on all DCs without additional updates.

Windows presets the RID Block Size registry subkey value to 0, which means that the internal default of 500 is used. You can't use the RID Block Size to set RID pool values lower than 500: It always defaults to 500. You can use it only to set higher RID pool values.

On machines running OSs earlier than Windows 2000 SP4, a flaw in the RID threshold compare logic caused RID Block Size values higher than 500 to revert back to the default allocation of 500. Win2K SP4 fixed this.

Windows 2000 DCs request a new RID when 20 percent of their RID pool remains. Starting with Win2K SP4, Microsoft increased the threshold at which DCs request a new RID pool to 50 percent. Therefore, a post–Win2K SP4 DC with a default pool size of 500 requests a new pool when 250 RIDs have been consumed.

To close, I want to share some practical thoughts on the domain-level RID pool that's controlled by the RID master and that can hold approximately 1 billion SIDs. If your domain were ever to reach the 1-billion RID limit, it wouldn't be able to create new user, group, or computer accounts.

Obviously, there's very little chance that any AD installation would ever reach this limit. Still, it's good to make sure that you don't have provisioning systems or scripts that automatically or accidentally bulk-create user, group, or computer accounts and that could suddenly eat a large piece of your RID pie.

To give you some peace of mind, you can check how many RIDs your RID master has already issued by using the Dcdiag command-line tool that's available on every Windows 2008 AD DC.

In Windows 2003, the Dcdiag tool is included in the Support Tools that are available from the product CD-ROM. In Windows 2000, Dcdiag is part of the Resource Kit. You can also download it from the Microsoft Download Center. To check the RID allocation with Dcdiag, type the following at a command prompt:

```
dcdiag.exe /test:ridmanager /v
```

Figure 1 shows the results of the command; the RID allocation appears in the RidManager section.

—Jan De Clercq

InstantDoc ID 141010

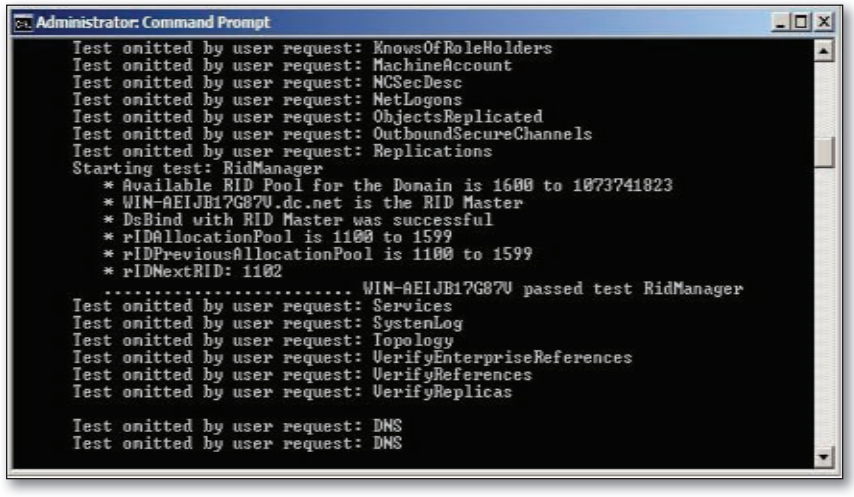


Figure 1: Using Dcdiag to see the RID pool allocation

## Q: Will the final version of System Center Virtual Machine Manager 2012 support Windows 8 Hyper-V?

**A:** No. System Center Virtual Machine Manager (SCVMM) 2012 will be released before Windows 8, so it won't have support for Windows 8 at release. However, after Windows 8 ships, expect an update to SCVMM to add support for Hyper-V hosts in Windows 8.

—John Savill

InstantDoc ID 141168

## Q: How can I establish recurring meetings with variations in Microsoft Outlook?

**A:** You might have a recurring meeting, such as a weekly departmental update, where there's a set resource such as a conference room that you use each time, but you need to make an exception. The exception could be an unusual, one-time occurrence, or perhaps a regular variation of the weekly meeting based on other criteria. An example might be that your usual resource room isn't available for your meeting on the first week of each month.

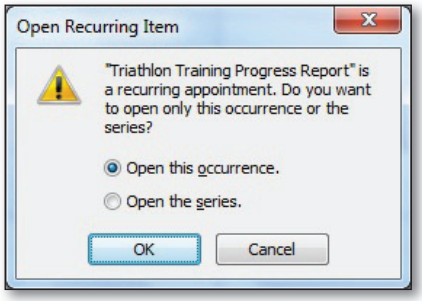


Figure 2: Making a change to a single instance of a meeting

There's no simple mechanism within Outlook to manage exceptions to recurring meetings. There are different ways to go about configuring the meetings within Outlook, however.

If you have a regular recurring meeting where every second meeting has a different feature—for instance, the meeting room location alternates—the best way to manage this variance is to submit separate recurring appointments to the participants, one for each resource. If you have a one-instance exception to a recurring meeting, then you can open the specific occurrence of that recurring meeting, make the necessary amendments, and resubmit it to the meeting attendees. For example, perhaps you have a set recurring meeting, but this week you have to change the location because of unforeseen circumstances. Find the meeting within the calendar and double-click to open it. Outlook prompts you to answer whether you want this single occurrence or all the remaining recurrences to view or edit. If you select *Open this occurrence,* as Figure 2 shows, you can make changes to the one meeting, rather than the series, and send an update to attendees.

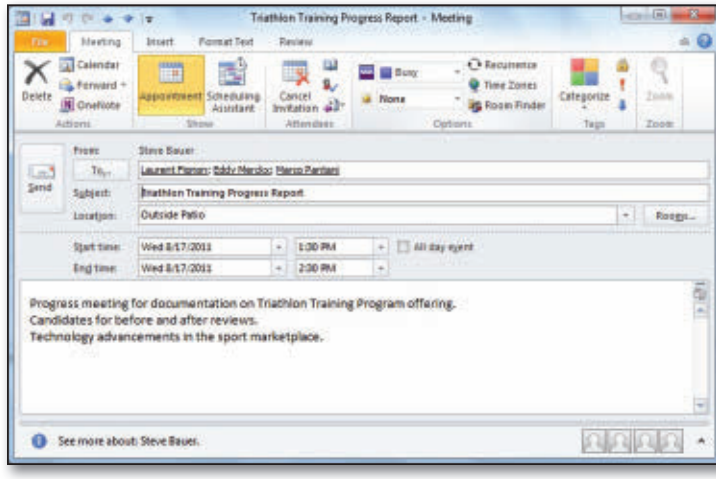


Figure 3: Creating a new meeting

When you create a new, recurring meeting with a variation of some sort, create the meeting with all the consistent components and save it before you make it a recurring meeting. Figure 3 shows a basic example of such a meeting.

In the Calendar View in Outlook, you can select the meeting you just created and use CTRL+C to save it to the clipboard and CTRL+V to save a copy of the item. You then open each calendar item, add the variation, such as different meeting rooms or different times of the day, set the recurring timing for each, as Figure 4 shows, and then send them as separate recurring meeting requests.

This method applies to any variation you might have in a recurring meeting, from a conference room, to time of day, or even specific attendees. Perhaps a certain attendee needs to attend only every fourth meeting, or the projector is needed only for the last meeting of each month. Whatever the difference, you can use a template meeting item with the common attributes of your recurring meeting and copy it, make the change, and save it as a second recurring meeting, complementing the original.

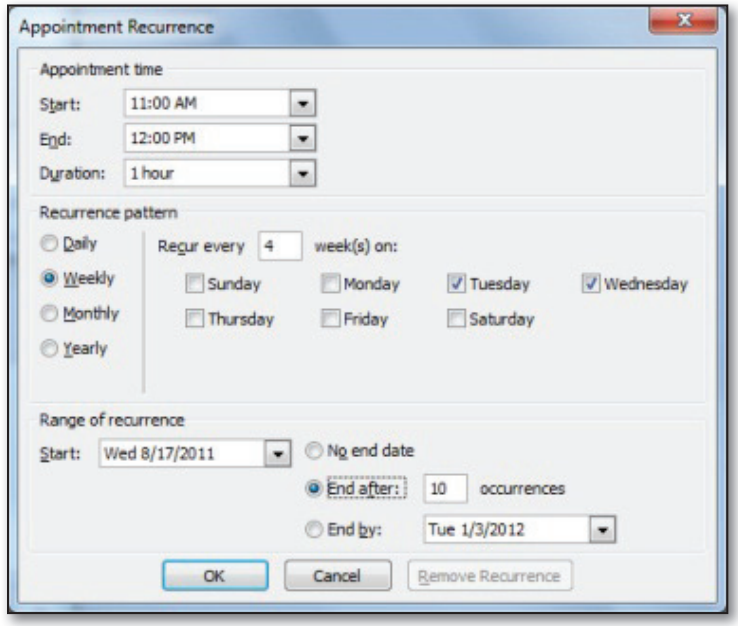


Figure 4: Setting the recurring timing for a meeting

Recurring meetings are valuable, timesaving tools for schedulers, but they can also be a hindrance when exceptions and changes are too frequent. Your fellow attendees might even tire of meeting updates. Getting meetings scheduled correctly the first time can be a big help.

—William Lefkovics

InstantDoc ID 140769

# BE THE FUTURE OF IT MANAGEMENT

## NOW@MMS2012

Microsoft Management Summit

**APRIL 16–20 @THE VENETIAN**
Las Vegas, Nevada

### INNOVATION. EXPERTISE. COMMUNITY.

Learn how to solve the tough IT challenges of today and prepare for the future. Get a peek at soon-to-be-released server, cloud, client and device management technologies.

**Don't get wait listed – register today! mms-2012.com**

REGISTER BEFORE JAN 27 & **SAVE UP TO $275** USD OFF THE STANDARD REGISTRATION PRICE.*

**Visit mms-2012.com today to learn more.**

*Discount is shown in United States Dollars only and may vary.

Microsoft

# vSphere 5 Raises the Bar for Enterprise Virtualization

VMware provides unmatched scalability and infrastructure automation

by Michael Otey

Although Microsoft has made decent progress in the virtualization race, there's no doubt that the leader in enterprise virtualization is VMware. VMware has been in the virtualization business since the company released VMware Workstation 1.0 way back in 1999. In addition, VMware was the first to bring bare-metal hypervisor-based virtualization to the enterprise, with the release of VMware ESX Server in 2001. VMware has a sizable lead in the virtualization market; current market estimates show that the company holds about 75 percent of the enterprise virtualization market. In June 2011, VMware released vSphere 5—even further raising the bar for enterprise virtualization. In addition, vSphere 5 continues to provide the foundation for leveraging virtualization to build the private cloud (see the sidebar "vSphere and the Cloud").

## ESXi Becomes the Hypervisor Standard

The core of vSphere has traditionally been VMware's ESX Server hypervisor. Eventually, VMware introduced the smaller-footprint ESXi hypervisor, which the company provided as a free product. The ESXi hypervisor essentially had the same virtualization capabilities as the larger ESX Server, but it jettisoned the service console—replacing it with a limited character-based UI and remote management through the VMware vSphere Client. The lighter-weight ESXi hypervisor also removes ESX Server's built-in web server, which means you now have to manually download the vSphere Client.

In vSphere 5, ESXi became the preferred hypervisor in part because of its smaller size but also because it has a reduced attack vector, making it more secure. In addition, the new ESXi for vSphere 5 features a built-in firewall that lets you limit traffic by IP address and subnet. vSphere 5 continues to support the older ESX Server hypervisor. The free version of the ESXi hypervisor is now called the VMware Hypervisor.

## Streamlined vSphere Editions

VMware streamlined the editions of vSphere 5 by eliminating the Advanced edition. For vSphere 5, VMware provides the Standard, Enterprise, and Enterprise Plus editions. The migration path from vSphere 4.1 Advanced is to vSphere 5 Enterprise. Table 1 summarizes vSphere 5's editions and their main features. All vSphere 5 editions use the ESXi hypervisor, and they all support vMotion, which provides the capability to move running virtual machines (VMs) between ESX and ESXi servers with no end-user downtime. Processor and vRAM entitlement are related to changes in the vSphere licensing scheme.

vSphere Standard is the starting point for most medium-sized business. This edition includes support for VMs with up to eight virtual CPUs, as well as support for VMware high-availability clusters and disaster recovery. The Enterprise edition is geared for larger businesses. It adds support for hot-add RAM and CPU to running VMs, as well as support for Storage vMotion and the Distributed Resource Scheduler (DRS). Storage vMotion lets a VM's files be moved to different storage locations with no downtime, whereas DRS

provides dynamic load balancing and power management for multiple ESX and ESXi hosts. The Enterprise Plus edition is the top of the vSphere product line. It provides all the features in the Enterprise edition, plus it adds the Distributed Switch capabilities, Host Profiles, and the new Policy Driven Storage support.

For centralized management, most organizations will also want to deploy VMware vCenter Server. vCenter Server is required to enable many of vSphere's advanced capabilities. It lets you provision, monitor, and manage VMs. Two editions of vCenter Server are available for purchase: vCenter Server Foundation and vCenter Server Standard. The vCenter Server Foundation edition lists for $1,495 and includes support for up to three vSphere hosts. The vCenter Server Standard edition lists for $4,995 and isn't limited as to the number of hosts it can manage. This edition also adds process automation through the VMware vCenter Orchestrator and multi-server insight with its vCenter Server Linked Mode.

## Support for Huge VMs

One of the most important changes in vSphere 5 is its enhanced support for highly scalable VMs. VMware has always been the leader in VM scalability, and with vSphere 5 the company extends its lead. vSphere 5 VMs can have up to 1TB of virtual RAM, which is four times the virtual RAM supported by any previous release. The Enterprise Plus edition also supports guest VMs with up to 32 virtual CPUs.

## vSphere Storage Appliance

One of the things that has slowed vSphere's adoption in small-to-midsized businesses (SMBs) is that you must have a SAN to take advantage of vSphere's high-availability features, such as vMotion. Many SMBs use DAS and can't afford a SAN. vSphere 5 introduces the new vSphere Storage Appliance (VSA) feature, which enables the creation of shared storage using DAS from two or three local ESX servers. Although its name implies that it's a hardware device, the VSA is actually a software feature. VSA divides the non-boot storage space of a vSphere server in half and makes each server node a primary for one volume and a replica for a second volume. The storage is replicated and half of the storage is assigned to the replica. If a node fails, an election occurs between the remaining nodes and the cluster directs the failover of the data store to the secondary replica.

## Policy Driven Storage

vSphere 5 Enterprise Plus offers support for Policy Driven Storage. Policy Driven Storage lets administrators create policies that determine where VMs should be located and moved using Storage vMotion. Storage policies associate VMs, data stores, and SAN devices with storage profiles. They're designed to ensure that VMs always run on the appropriate storage location to meet service level agreements (SLAs) for performance and to accommodate the VM's space requirements. A storage profile identifies the storage characteristics that a particular VM should have.

When a VM is created, it can optionally be associated with a storage profile. The storage policy determines the appropriate storage locations that meet the profile's requirements, and the administrator is prompted for the desired appropriate location. Likewise, if a Storage vMotion process is initiated, only target storage locations that meet the profile's storage characteristics will be used.

## Storage DRS

Another storage-related enhancement in vSphere 5 Enterprise Plus is the addition of Storage DRS. vSphere uses Storage DRS to balance the placement of VM files

# vSphere and the Cloud

Selling the cloud, the private cloud, and the hybrid cloud seems to be the top goal of every major IT vendor these days—and VMware is no exception. VMware was actually one of the leaders in the cloud space to show how virtualization can act as the foundation for cloud computing. Not surprisingly, since VMware has no global infrastructure services to sell, VMware places most of its emphasis on the private cloud and the hybrid cloud.

The concept of the public cloud is reasonably clear. A vendor provides a set of services that a customer can subscribe to. These services are typically Infrastructure as a Service (IaaS), Platform as a Service (PaaS), or Software as a Service (SaaS). IaaS typically means you're leasing VMs that are hosted on an Internet vendor's infrastructure. An example of IaaS is Amazon's Elastic Compute Cloud (EC2). PaaS typically means you lease operating services from a cloud vendor. Windows Azure is an example of this. With SaaS, you lease an application provided by a vendor. Microsoft Office 365 and Salesforce.com are examples of SaaS. Costs are typically metered by usage. Advantages of the cloud include lower capital expenditures and operating costs, as well as increased flexibility and scalability.

The private cloud is a newer concept and not quite as well understood. The main idea behind the private cloud is to take your existing infrastructure and make it more flexible, dynamic, and automated. Virtualization and technologies such as vMotion, Storage vMotion, and the Distributed Resource Scheduler (DRS) make this possible by automatically migrating workloads between pools of infrastructure. The primary advantage of the private cloud is that it moves your infrastructure away from being fixed, where one resource handles one workload, to a more fluid and dynamic infrastructure, where workloads can be automatically matched and moved to the resources that can best support them. The private cloud can also improve operational efficiency. In periods of low utilization, unused resources can be shut down. VMware's DRS performs both of these functions. The private cloud can also lay the groundwork for chargeback computing, in which IT consumers can be changed based on their resource utilization. VMware vCenter Chargeback performs this function.

VMware's principal cloud management product is vCloud Director. vCloud Director lets you provide and manage IaaS across multiple clusters in your data center. VCloud Director enables the creation of virtual data centers, as well as rapid provisioning of virtual machines (VMs) and virtual applications (vApps).

InstantDoc ID 141318

Table 1: vSphere 5 Editions and Specifications

| Specifications | vSphere Standard | vSphere Enterprise | vSphere Enterprise Plus |
|---|---|---|---|
| Price | $995 | $2,875 | $3,495 |
| Processor Entitlement | Per 1 CPU | Per 1 CPU | Per 1 CPU |
| vRAM Entitlement | 32GB | 64GB | 96GB |
| Max Host RAM | 256GB | 384GB | 576GB |
| Max VM CPUs | 8 | 8 | 32 |
| vMotion | Yes | Yes | Yes |
| HA & DR | Yes | Yes | Yes |
| Hot-Add RAM/CPU | No | Yes | Yes |
| vShield Zones | No | Yes | Yes |
| Fault Tolerance | No | Yes | Yes |
| Storage vMotion | No | Yes | Yes |
| DRS | No | Yes | Yes |
| Distributed Switch | No | No | Yes |
| Host Profiles | No | No | Yes |
| Policy Driven Storage | No | No | Yes |
| Storage DRS | No | No | Yes |

across data stores based on I/O metrics and data store capabilities. To accomplish this, Storage DRS captures latency information regarding all the data stores in the cluster. If the latency time for a given data store is above the threshold value for a significant percentage of time over a specified period, then Storage DRS automatically invokes one or more Storage vMotion operations to rebalance the VMs in the data store cluster.

## vCenter Server Linux Virtual Appliance

vSphere has always required vCenter Server to provide centralized management, provisioning, updating, and load balancing of VMware vSphere hosts. In all previous releases of vSphere, vCenter Server only supported the Windows Server OSs. In vSphere 5, vCenter Server can run from the vCenter Server Appliance (vCSA).

The vCSA is a prebuilt VM that runs SUSE Linux Enterprise Server 11 and a new Linux version of vCenter Server. The vCSA comes with an embedded version of IBM's DB2 and can support up to 50 VMs. For greater scalability, it can be configured to connect to a separate DB2 or Oracle instance. The vCSA doesn't support Microsoft SQL Server. The vCSA is initially configured using a web browser; however, after the initial configuration is complete, you can use the vSphere Client to manage it exactly like the Windows version. The vCSA supports authentication using Active Directory (AD) or Network Information Service (NIS). Notably, the current version doesn't support IPv6.

## New vSphere Essentials

VMware and vSphere are clearly entrenched in the enterprise—an approach that overlooks SMBs. Traditionally, VMware and vSphere have been too expensive and too complex. With vSphere 5, VMware aims its new vSphere Essentials Kits at the lower end of the market. VMware offers two versions of the Essentials Kit: the VMware vSphere Essentials Kit and the VMware vSphere Essentials Plus Kit.

Both vSphere Essentials Kits provide support for six CPU entitlements, 32GB of vRAM, and up to eight virtual CPUs per guest. The main difference between the two editions is that the VMware vSphere Essentials Plus Kit provides support for the High Availability, Data Recovery, and vMotion features that aren't available in the less-expensive VMware vSphere Essentials Kit. The VMware vSphere Essentials Kit lists at $495 and requires a one-year support package. The VMware vSphere Essentials Plus Kit has a list price of $4,496. For centralized management, vCenter Server for Essentials is integrated into both vSphere Essentials Kits.

## Licensing Changes

Although not a technology feature, one of the most significant changes VMware introduced with vSphere 5 was a revamped licensing model. VMware has always been known to be expensive. With vSphere 5, VMware began licensing vSphere according to a more cloud-like, pay-for-consumption model.

The new vSphere 5 licensing model is still based on processor licenses, but it removes the physical restrictions of CPU cores and RAM per server. Instead, it bases the licensing on pooled virtual memory (vRAM) usage (where vRAM is defined as the memory configured to a VM). Each physical processor CPU in a server needs to have a vSphere 5 processor license. The vRAM usage is calculated using a 365-day moving average of the high daily watermark of vRAM for all powered-on VMs. There's no automatic shutdown if the licensed limits are exceeded. However, vCenter issues an alert if the available pooled vRAM is exceeded.

Customers were quite resistant to this licensing change initially, because it resulted in higher costs—especially in situations in which memory was overcommitted. This change made high levels of server consolidation (traditionally a very desirable factor) more expensive than in the past. In addition, the new model could penalize customers for overcommitting RAM.

VMware quickly revised the vSphere 5 licensing scheme, essentially doubling the vRAM limits so that most existing installations wouldn't need to upgrade their licenses. Although large businesses will typically take these licensing changes in stride, such changes could certainly push small customers to either hold on to vSphere 4.1 or to seriously consider Microsoft's less-expensive Hyper-V product.

## vSummary

VMware's vSphere 5 continues to raise the bar for enterprise virtualization. Although Microsoft Hyper-V is gaining traction at the low end of the market, vSphere 5 remains the undisputed leader in the enterprise space. vSphere 5 offers unmatched scalability and infrastructure automation. They say there's no such thing as a free lunch, and vSphere 5 is a good example of that. vSphere 5 offers a premium virtualization feature set, but at a premium price.

InstantDoc ID 141317

**Michael Otey**
(motey@windowsitpro.com) is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).

FEATURE

# Jeffrey Snover Discusses Windows Server 8

## Microsoft's lead architect for the Windows Server Division on the latest server operating system

by Sean Deuby and Michael Otey

Jeffrey Snover is lead architect for the Windows Server Division at Microsoft. He has been in the IT industry for over 30 years and with Microsoft since 1999, and he invented Windows PowerShell. *Windows IT Pro* technical directors Sean Deuby and Michael Otey sat down with Snover at the Windows Connections conference in early November to discuss some of the latest enhancements to Windows Server.

**Sean Deuby:** What does "lead architect for Windows Server" mean? What does that role look like on a day-to-day basis?

**Jeffrey Snover:** It's a couple of things. First, we have people who are very specific to a feature—like the architect for that feature. My role is to look across the scenarios to make sure they're all fitting together. And then in general, the job of an architect is to be what I call the "guardian of the long-term," which is to say, "Hey, that's great, but how is that going to fit in the next release and the release after that?" So that's what I do for Windows Server—figure out not just how to deliver a great product in this release, but also use this release to set up the next one and the one after that.

**Michael Otey:** Sean and I attended the Windows Server 8 Reviewer's Workshop in September, and we were blown away by all the new features coming up in Server 8. I know there are way too many to talk about, but what are some of the highlights for you of what we're going to see?

**Snover:** As a technologist, I look at the technology innovations. Although we've had great technology innovations in the past, I think this is by far the largest, most transformative release we've ever had. We have major innovations in storage. Honestly, in the past, we've had some weakness in our storage stack and people had to buy very expensive, very high-end storage arrays to do some of the things they wanted to do. Now if you have those things you're going to get more value out of them because we have a close partnership with those storage vendors. A lot of the things you think you could only get from the storage arrays, you're going to get with in-box storage. If you sit down and look at the details, in every single layer of the storage stack there's transformation—the way we deal with disks, the way we deal with the file system, the way we cluster things together.

**Otey:** Some of the things that really jumped out at me from the storage side were the built-in data deduplication capabilities, which are pretty amazing; the total revamp of the Checkdisk operations, which are much more efficient and online and dynamic; and the ability to take advantage of the storage back-end arrays, where you wouldn't have to funnel the I/O through the servers and instead when you're doing a file-copy operation on that back-end array you can tell it to take advantage of those kinds of things that are built in. Those are big changes, and they're obviously baked into the hardware at a pretty deep level and into the OS.

**Snover:** Right, and in the whole storage space—the ability to take a bunch of inexpensive disks and pool them together and do thin provisioning and . . . this delta. In the past, NTFS was designed around SCSI. But it turns out, some of the inexpensive SATA drives didn't correctly implement a number of the commands. They looked good on a benchmark, but not so good when there's a power outage and now we don't have [any data]. Now we detect that and modify our flushing algorithms to ensure consistency and get great reliability, from notebook drives all the way up to very large sets of SATA

drives—so now you can more safely take advantage of commodity components.

**Deuby:** That raises an interesting point. A lot of companies are still upgrading to Windows Server 2008 R2, even though Server 8 will be out soon. There's a big push toward private cloud, and a lot of people are wondering how to manage the private cloud. Storage is one of the reasons you should migrate to Server 8 rather than stick with Server 2008 R2 as you're building your private cloud and your next-generation infrastructure.

**Otey:** Another thing we were really impressed with is some of the changes in the new hypervisor and virtualization of Server 8. Can you tell us about those?

**Snover:** First is scale, scale, scale. That's not just limited to virtualization. There's been a strong push from the very beginning on scale—finding out, throughout the stack, where the bottlenecks are and fixing them. So now we go up to 640 CPUs. It's phenomenal. When you have a virtualized machine, you can now have 160 processors and 2TB of RAM, and then the VMs themselves, 32-processor VMs and 512GB of RAM.

**Deuby:** This is version 3 of the hypervisor; can you tell us about v3?

**Snover:** In the vast majority of my 31 years in the industry, I worked for companies that competed against Microsoft. Microsoft would enter a market and would quickly point out the competition's failings and flaws. We'd pat ourselves on the back and feel very confident—but honestly, we always knew: Sell your stock options before Microsoft's version 3 hits the ground, because by version 3 they've dialed it in and figured it out. They do a great job with version 3. And that's certainly the case with hypervisor and virtualization. It's not just about scale, but also the ability to do replication—definitely the best replication story out there, achieving what some people call the Holy Grail of replication, where not only can I do replicas based on a cluster, I can do it without a cluster using shared storage, or just having an Ethernet cable.

And you have the ability to do that replication synchronously or asynchronously and the capability for disaster recovery scenarios, where you can say, "Hey, I'm running this here and asynchronously replicating it, perhaps in the cloud to a hoster in case anything goes wrong here." Machines go down—but sometimes entire sites go down, so the ability to inexpensively back up to the cloud is a wonderful thing.

**Deuby:** I think you bring up a key point, which is the inexpensive part of it. So, small-to-midsized businesses don't have to pay an arm and a leg.

**Otey:** That's true. VMware has been criticized for being an expensive solution, and it seems like it becomes more expensive all the time. If Hyper-V is built in to Server 8, it's a good value proposition for SMBs.

**Snover:** This is Microsoft's history and our distinct competence: the ability to take very high-end, very expensive computing and make it available to the masses. You see that with virtualization: very high-end and increasingly expensive. I think [VMware] became aware that v3 is coming out, so it's jacking up prices to get the money while it can. You see that with virtualization, you see it with storage, you see it with management.

Another example is remote direct memory access [RDMA]. It lets me say I've got a specialized NIC that allows me to have an alternate network path to TCP. It's amazingly fast, amazingly low-latency because it's all done in the hardware. In the past, that was really done by the high-performance computing world. So, *x* thousand guys pay through the nose to get these great NICs, get fantastic performance, but now what we're saying is that everybody should be delivering NICs like that because in addition to those scenarios, which continue to exist, we now have a kernel-mode API that can access that and we take advantage of that kernel-mode API with SMB [server message block]. So now SMB Direct gives the ability to use this RDMA and go as fast as the wind.

Microsoft's engineers changed the protocol so we can use multiple TCP connections on the same SMB session. This means you can have as many NICs as you want, connected between source and destination. And the session is dynamic, so you can remove a NIC and dynamically adjust. you can add a NIC and dynamically adjust. So you have maximum bandwidth and resiliency. This was the most impressive thing, in my opinion—the fast failover.

In the past, if you clustered your file server, we've raised the bar from high availability to continuous availability. High availability says that if something goes wrong, you can fail over and restart your operation and succeed. Continuous availability says that if a failure occurs, we detect it and resolve it quickly enough that the application never notices it. The operation

takes a little bit longer, but it doesn't time out—which is a lot of hard work.

**Otey:** One of the other points you touched on is that Server 8 now has built-in NIC teaming, so you don't need specialized vendor NICs to get this kind of availability.

**Snover:** You can say, "I got that from my vendor in the past." Well, yes and no. You could, but it only worked with that vendor's NIC. You couldn't have heterogeneous NICs. If you ever had a problem and you called Microsoft and said, "I'm using NIC teaming," Microsoft would say, "OK, turn that off—that might be the issue." But now we support it, so if you call, we'll help you through it. But 32 NICs—it's just phenomenal. The performance team did such a good job paying attention to the NUMA algorithm's uniformed architecture.

**Otey:** And that's especially important for performance in VMs.

**Snover:** Yes, because you can't buy a server today that's not NUMA capable. So according to NUMA, there are things that are cheap and there are things that are expensive to do, and the software has to be aware of that and pay attention to it—otherwise, you have bad performance. So we've gone through the entire stack looking for these problems. The receive-side scaling, which is to say, I've got a lot of bandwidth coming in and it all goes to the same processor. But you can only go so big—so you want to fan it out in a way that's aware of the NUMA topology. So you're not just saying, "I fan it out to all these nodes, but they're all in the same NUMA node"—because then that doesn't fan out. So it's just fantastic scaling.

**Otey:** So, that can give you linear scaling as you're moving up, as far as processing and cores and that type of thing.

**Snover:** Did you see those numbers? I had to go back and say, "You need to check your numbers, because I've never heard of this." The scalability of VMs and basically—it's 8 to 16—for a SQL Server workload. You went from an 8-CPU VM to a 16-CPU. I think it's 1.7x scaling, which is just phenomenal. You'd expect 1.4, 1.5, and let me shake your hand. But here's one, I didn't think it was possible, but from 16 to 32, it was 1.9—those are crazy numbers.

And it's not just a bolt-on. It requires work at every layer of the stack—and in the management space, the same thing. We have this new multi-machine Server Manager, but in fact that's just a very thin layer on top of the multi-machine management capabilities within the OS. And it changes at the protocol level, at the PowerShell level. They had to make changes to WMI. At each layer, we had to make changes to be able to support that.

**Otey:** You touched on something that's going to be super important with Server 8, which is the changing management paradigm. With Server 8, you've taken a different look at how admins should manage Windows Server.

**Snover:** Absolutely. In the past, you bought a server, and a full server was the default. You got a GUI with it, and there was Server Core and a few specialized people would use Server Core, but there were a lot of issues with it. So with each release we've invested in Server Core, made it better and better, made it able to support more roles, be able to do more manageability.

With Server 8, we're now confident enough to say that Server Core is the preferred management deployment role. Full Windows Server is still there as a compatibility mode, but by and large we want everybody to use Server Core, which is basically to say "headless server."

We still support GUIs—we're not walking away from GUIs. GUIs are what make the company great. GUIs help customers, but those GUIs should run on the client, and the client consumes as much CPU and as much memory as you want—it's a client. Obviously you don't want that on a server. And then layer that GUI on top of PowerShell, remote PowerShell, so that it can do multi-machine management, and anything I can do from the GUI, I can then automate.

**Otey:** So you're saying out of the box, Server Core is going to be the default installation option. But in the past, it was difficult to switch back and forth between Server Core and the full installation. Has that changed?

**Snover:** Yes. Why are we confident? Basically, three answers: In the past, you chose Server Core or Windows Server—and if you made a mistake, you started over. Now you can go from Server Core to full Server and back again. And there's something in between, which is to say that with full Server, you can take off the Metro shell in IE. So you can still run GUIs, you can still run Server Manager, but you launch it from the command line. This gives you many of the benefits of Server Core in terms of reduced footprint and reduced serviceability, which means it takes fewer patches. For those people who've been able to make the full transition—maybe the admin hasn't been able to do that, they're not fully cognizant of PowerShell, or remote management, or an application. Often what we've found in our compatibility tests is that an application will require the GUI for installation but not operation.

I mentioned three things—that was one. Now it's safer; reduce the risk. The second is manageability. In the past, PowerShell 1.0 shipped with about 130 cmdlets; PowerShell 2.0 shipped with 230 cmdlets; and now we ship with more than 2,300 cmdlets—so over a factor of 100 more cmdlets.

And now, you can really do full management of the box locally. And if you want to, we now have remote management. In the past, Server Manager couldn't remotely install a role. But now you can. You still use the GUI, but you do it remotely.

The third thing is role availability. There were certain roles that required full Server. Now, more and more of those roles require Server Core. And more importantly, the Denali release of SQL Server [SQL Server 2012] runs on Server Core. So we're feeling pretty confident. This is certainly one of the strong messages we have for everyone in the community, for the ISVs: Love the GUI, just don't run it on the server. Run it on the client, and use PowerShell remoting to the server.

**Deuby:** Much of what Server 8 is focused on is helping customers build their own private cloud—and certainly it will be used as a major component of the public cloud as well. Are there any enhancements that have been made to identity to help with the integration because we're looking at building private cloud now and going to something that's hybrid? So the ability to

have some portability between the two certainly has something to do with Active Directory Federation Services [AD FS]. Has anything been done in that area?

**Snover:** The big investment there is in the area of roles-based administration. A lot of this isn't just the technology itself. I recently gave a talk in Japan, during which I was trying to explain our investment in continuous availability. I said that basically we're trying to treat three things: more 9s, more 9s per dollar, and easier 9s. Which is to say that you can have all this capability, but it requires regular people to be able to put it together. If you require the world's best admins—at their best, fully trained—then you're not going to get that continuous availability. It's the same sort of thing in the area of identity and access control. We have these great capabilities, but sometimes they can get like, "What? Do I understand this?" One of the big things we've done in Server 8 is take the roles-based access and . . . .

**Deuby:** You're talking about flexible access control.

**Snover:** Exactly. The ability to say, "Hey, these roles—I want to automatically identify these files with a set of attributes, high business impact, who owns it, et cetera," and then say, "Hey, you can have read access if you belong to this department." So it's not just identity and group but a richer set of access rights.

**Otey:** Windows 8 and Server 8 share the same kernel, right?

**Snover:** Yes, always. And that's why—same kernel, same GUI, there's one Windows. It takes on different flavors, but there's one Windows. Some people ask, "Why does Windows Server adopt the Metro UI?" I don't understand the question. There's only one Windows. So if Windows has a new UI, Windows has a new UI. And I get why you might not want that on your server, which is why we have Server Core. But it's actually quite a great UI on the client desktop, especially if you have Touch. And they've done some great stuff, it's just that you wouldn't want that consuming resources of your server. One thing I try to point out is that if you have SQL Server, you want every single transistor delivering transactions, because that's its job. When it comes to GUIs, guess what? We want to deliver. I'll consume all available CPU cycles, memory, to deliver the experience because that's all about you.

**Otey:** A lot of administrators aren't familiar with PowerShell. They've struggled to get into it; they kind of know Windows Shell scripting, maybe they know VB Script, but PowerShell is a bit daunting to them. What kind of advice can you give those admins to help them adopt PowerShell, to help them move toward it and take advantage of it?

**Snover:** PowerShell is the glue coat. We've glued things together. So we deal with the world as it is. It's a messy world. Ultimately, we're trying to drive to these cmdlets—these high-level task-oriented abstractions that allow people to think about what they want, type it, and get it. Ultimately, perfection would be do/myjob/ordermeapizza. Obviously we're not going to get there. But if you think about what you want, and you can type it and get it, PowerShell is very easy. The fact that you have to type it isn't a big issue. It's a different input device, but that's not the issue. You think about something; you type it and get it. In the past, you had to do some COM programming or find some WMI classes, or invoke some command-line shell and parse the output—which can get pretty rough. Some people just love that stuff and are very successful at it. But ultimately a lot of people just want to type it and get it. So this is where the 130, 230, 2,300 cmdlets come in.

**Deuby:** I know there are features—for example, in Active Directory I think it's called a cmdlet window—where you can see commands going past as you're manipulating users and things like that. You can see what they're doing, kind of giving you feedback to it. As we're moving to various degrees of comfort toward a cloud computing era, one of the key tenets is automation. And if you're going to have any kind of an automation that involves Microsoft products—for the IT pros out there, you have to learn PowerShell.

**Snover:** We're working on something—which will ship before Windows 8, and ship out of box—that's a script-sharing facility, which allows people to say, "Hey, I'm interested in something around virtualization or Exchange," and it has this search aggregation service that looks at all these various script repositories, aggregates them, and delivers you the results, along with the reputation service. So you can say, "That looks good," and then you can copy the script and make it your own. It's very clever technology. In fact, you can configure it so that you can set up a local script-sharing repository as well. So a big company can say, "Point this thing at my script repository—I want you using ours, not the community's."

There are some big challenges for the IT pro community. The cloud is wonderful, but it requires changes on the part of IT pros. A bunch of people are going to prosper through this change; others will not. It's all about automation. I like to joke, but it's true: Bill Gates will pay any of us a million dollars a year if we can produce enough value. The trick is, how do you produce enough value? If you just take a GUI and a mouse and go at it, it's hard to differentiate a really smart guy from a not-so-smart guy with a mouse. Can you click faster? Automation is basically a skill amplifier. You're able to be much more productive. That's the test of a technology—to put it in front of a junior person and in front of an expert person and see how much more of a delta you can be. PowerShell acts like a big amplifier. People who are skilled can produce a lot more—a ton more productivity.

InstantDoc ID 141366

**Sean Deuby**
(sean@windowsitpro.com) is technical director for *Windows IT Pro and SQL Server Pro*, and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.

**Michael Otey**
(motey@windowsitpro.com) is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).

# Use LDAP over SSL to Lock Down Active Directory Traffic

Learn how to configure LDAPS on your AD domain controllers

by Jan De Clercq

The standard protocol for reading data from and writing data to Active Directory (AD) domain controllers (DCs) is LDAP. AD LDAP traffic is unsecured by default, which makes it possible to use network-monitoring software to view the LDAP traffic between clients and DCs. This security problem also applies to the LDAP subprotocols, such as LDAP bind, that applications, services, or users use to transport credentials and authenticate against a Windows DC.

Organizational security policies typically require that all client/server communication is encrypted. In addition, applications that integrate with AD might require encrypted LDAP communication.

To make LDAP traffic secure, you can use the Secure Sockets Layer/Transport Layer Security (SSL/TLS) protocols; this combination is referred to as LDAP over SSL—or LDAPS. To ensure that no one else can read the traffic, SSL/TLS establishes an encrypted tunnel between an LDAP client and a Windows DC. In this article, I explain how to set up LDAPS on the DCs in your Windows Server 2008 AD infrastructure.

## LDAPS Server Certificate Requirements

LDAPS requires a properly formatted X.509 certificate on all your Windows DCs. This certificate lets a DC's LDAP service listen for and automatically accept SSL connections for both LDAP and Global Catalog (GC) traffic. The server certificate is used for authenticating the DC to the client during the LDAPS setup and for enabling the SSL communication tunnel between the client and the server after setup. As an option, you can use LDAPS for client authentication—but doing so requires that you also install a client authentication certificate on each of your clients.

In the next section, I explain in detail how you can obtain an LDAPS server certificate for your DCs—but first, let's look at what rules this certificate should adhere to.

The LDAPS certificate and its associated private key must be stored in the DC's Personal certificate store (also referred to as the MY certificate store). To view the content of your DC's certificate store, follow these steps:

1. On the DC, click Start, type mmc, and click OK.
2. Click the File menu option, then click *Add/Remove Snap-in.*
3. Click Certificates, Add.
4. In the Microsoft Management Console (MMC) Certificates snap-in dialog box, select *Computer account* and click Next.
5. In Select Computer, select *Local computer* and click Finish.
6. In *Add or Remove Snap-ins,* click OK.
7. In the console tree, expand Certificates (Local Computer), then the Personal container, and finally the Certificates container.

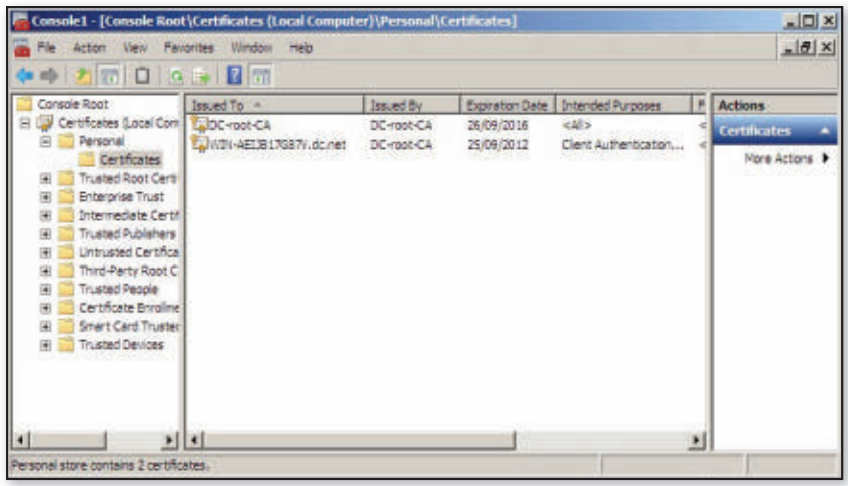


Figure 1: Certificates in a DC's Personal certificate store

8. In the right-hand pane of the Certificates snap-in you'll see a list of all the certificates that are stored in your DC's Personal certificate store, as Figure 1 shows.

The LDAPS certificate must meet the following X.509 certificate extension requirements:

- The Extended Key Usage certificate extension must include the Server Authentication Object Identifier (OID): 1.3.6.1.5.5.7.3.1.
- The AD Fully Qualified Domain Name (FQDN) of the DC (e.g., mydomaincontroller.company.net) must appear in either the Common Name (CN) of the Subject field certificate extension or in the DNS entry of the Subject Alternative Name (SAN) certificate extension.

You can easily check the content of a certificate's X.509 extensions from the Windows Certificate Viewer, which you can also access from the Certificates snap-in. Double-click a certificate in the right-hand pane of the Certificates snap-in, then click the Details tab, as Figure 2 shows.

As for any certificate, the LDAPS certificate must have been issued by a Certification Authority (CA) that the DC and the LDAPS clients trust. Trust is established by configuring the clients and the server to trust the issuing CA's certificate (in a one-tier CA setup) or the certificate of the root CA to which the issuing CA of the LDAPS certificate chains (in a multi-tier CA setup).

You can find a list of all trusted CA certificates in the Trusted Root Certification Authorities container of a machine's certificate store. This store contains the certificates of CAs that you or your domain administrator consider trustworthy and that Windows can use as a trust anchor for validating other certificates. The CA certificates of CAs installed on machines that are part of your AD infrastructure are automatically added to client machines' certificate stores through the Group Policy Object (GPO) update mechanism. If the CA certificate of the CA that issued the LDAPS certificate isn't present in the Trusted Root CA Certification Authorities container, you can manually import it using the Import option that's available from this container's context menu.

Another important condition is that not only the LDAPS certificate but also a private key that matches the certificate is present in the DC's certificate store. To check whether the DC certificate store holds a matching private key for a given LDAPS certificate, you can use the General tab of the Certificate Viewer, as Figure 3 shows. If the correct private key is present, the bottom of this dialog box should show a key symbol and the text *You have a private key that corresponds to this certificate.*

The private key should also not have strong private key protection enabled—which means that Windows shouldn't prompt for a password each time the key is accessed. Strong private key protection is never enabled by default in Windows, which also applies to LDAPS certificates.

To validate whether your DC has a valid certificate from the command line, you can use the certutil utility as follows:

```
certutil -verifyStore MY
```

For more information about how to leverage the output of this command for troubleshooting LDAPS certificates, see the Microsoft TechNet article "Troubleshooting LDAP Over SSL" at blogs.technet.com/b/askds/archive/2008/03/13/troubleshooting-ldap-over-ssl.aspx.

## Obtaining a Server LDAPS Certificate

Your DCs will automatically receive a valid LDAPS certificate when you install an enterprise root CA (i.e., an AD-integrated CA) on one of your domain servers or DCs. However, this certainly isn't the most secure or most flexible option for providing

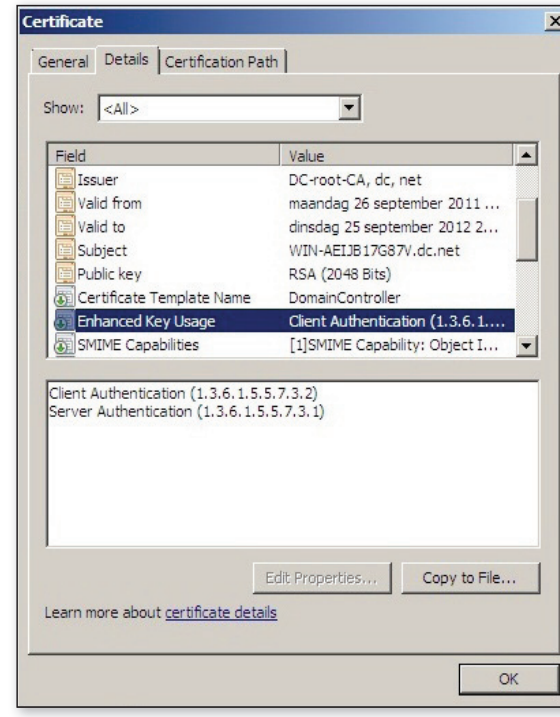


Figure 2: Certificate extensions in the Details tab of the Certificate Viewer

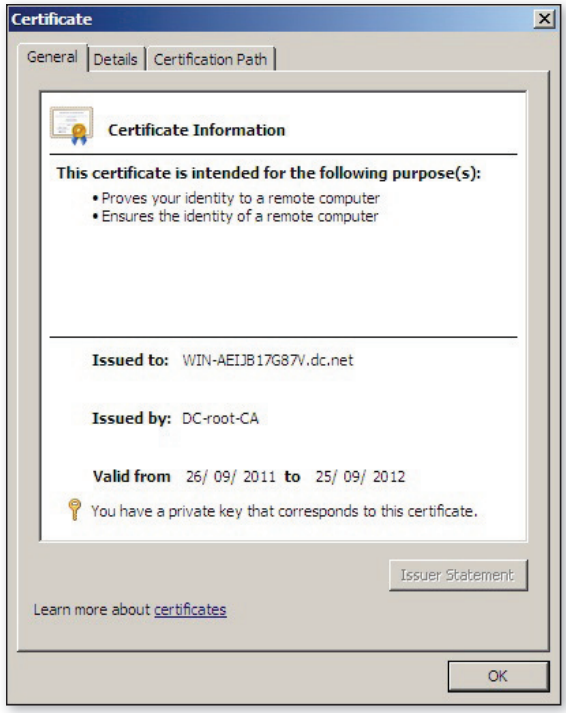


Figure 3: Private key properties in the General tab of the Certificate Viewer

certificate services to your users. Most organizations that want to bring a higher level of security and flexibility to their certificate services opt for a multi-tier Windows CA hierarchy that consists of a root CA and different issuing CAs. If your organization has created or plans to create a Windows CA hierarchy, I advise you to follow my instructions in this section for creating LDAPS certificates for your DCs. LDAPS certificates can also be issued by a non-Microsoft CA. (For detailed instructions, see the Microsoft article "How to enable LDAP over SSL with a third-party certification authority" at support.microsoft.com/kb/321051.)

When you have a multi-tier Windows public key infrastructure (PKI) hierarchy, you must first create a custom certificate template for LDAPS certificates in AD, then enable this custom template on all your issuing CAs, and finally manually enroll all DCs for an LDAPS certificate that's based on this custom template.

To create a custom certificate template for LDAPS certificates in AD, open the MMC Certificate Templates snap-in on one of your enterprise (AD-integrated) CAs. (Click Start, type mmc, and click OK. Click the File menu option, then click *Add/Remove Snap-in*. Click Certificate Templates, Add, OK.) In the Certificate Templates snap-in, expand Certificate Templates, right-click a template in the right-hand pane (e.g., the Kerberos Authentication template), and select Duplicate Template. Note that you can also duplicate another template (e.g., the Domain Controller Authentication template) as long as the template has the Server Authentication OID in its Extended Key Usage certificate extension.

In the Duplicate Template dialog box, leave the default Windows Server 2003 Enterprise option selected and click OK. This will make the properties of the new template appear, as Figure 4 shows. You should pay special attention to the following properties of the new template:

- On the General tab: Enter a template display name (e.g., "LDAPS"), set the validity and renewal periods (ensure that they're set according to your organization's certificate policy), and specify whether you want to publish the certificate in AD (select the *Publish certificate in Active Directory* check box).
- On the Request Handling tab: Ensure that the minimum key size is set according to your organization's certificate policy, and select whether the private key must be marked as exportable (select the *Allow private key to be exported* check box). You must mark the private key as exportable if you want to import the certificate into the AD NTDS certificate store, as I explain later.
- On the Subject Name tab: Ensure that the DNS name and Service Principal Name (SPN) are selected.

Finally, click OK to close the template properties and complete the new template customization.

To enable your issuing CAs to issue certificates based on the new LDAPS template, you must add the new template to the CA's Certificate Templates container. To do so, start the MMC Certification Authority snap-in on one of your enterprise CAs, expand the CA container, right-click the Certificate Templates container, and select New, *Certificate Template to Issue*. In the Enable Certificate Templates dialog box, which Figure 5 shows, you can then select the name of the newly created template. To close the dialog box, click OK.

As a last step, you must request LDAPS certificates for each of the DCs that requires LDAPS connections. To do so, perform the following steps on all affected DCs.

1. From the MMC Certificates snap-in, open the Personal container of the machine's certificates store, as I explained in the previous section on the LDAPS server certificate requirements.
2. Right-click the Certificates container and select All Tasks, Request New

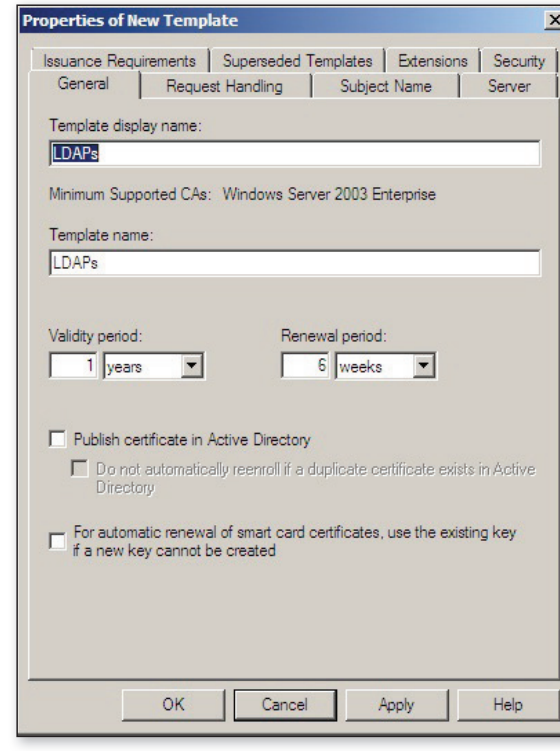


Figure 4: Certificate Templates properties

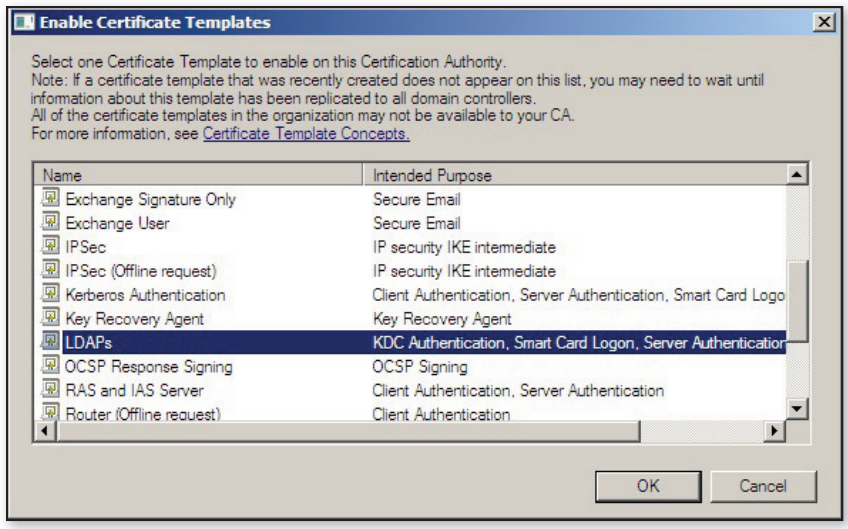


Figure 5: The Enable Certificate Templates dialog box

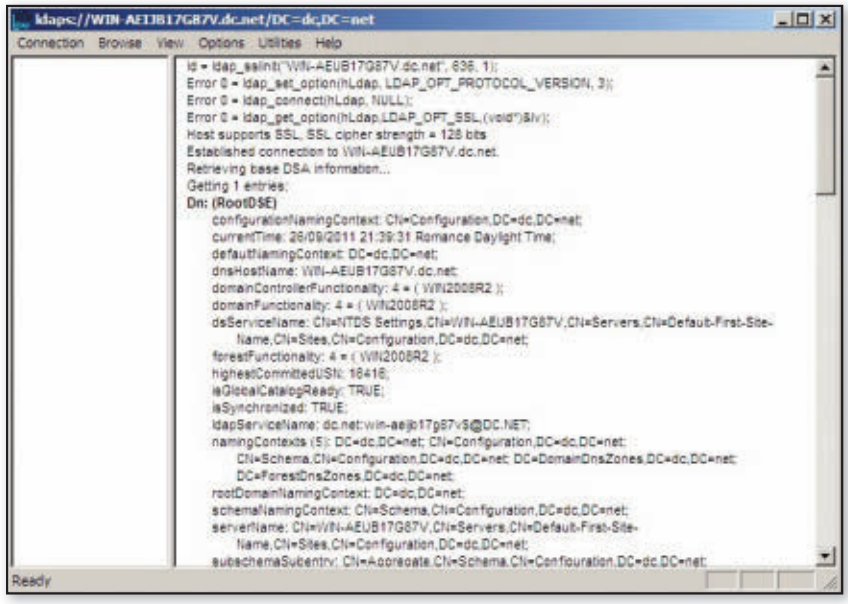


Figure 6: Verifying LDAPS connectivity using LDP

Certificate. This action will launch the Certificate Enrollment wizard. Click Next.

3. On the Select Certificate Enrollment Policy page of the wizard, leave the default of Active Directory Enrollment Policy and click Next.
4. Select the LDAPS certificate template and click Enroll.
5. Ensure that the enrollment succeeds and verify the properties of the new LDAPS certificates using the View Certificate option in the Details section.
6. Click Finish to close the wizard.

Windows Server 2008 provides a new option that lets you store the LDAPS certificate of a DC in AD's Personal certificate store on the DC. This is a good option if your DCs have multiple certificates with the Server Authentication OID in their Local Machines Personal store. In that case, it's difficult to predict which certificate AD will pick for LDAPS authentication. The new Windows Server 2008 logic makes AD first look for server authentication certificates in the AD certificate store. Therefore, this new feature can force AD to use the server authentication certificate that you generated using your custom LDAPS template. For more information about how to add the certificate to the AD service's Personal certificate store (also referred to as the NTDS certificate store), see the Microsoft TechNet article "Event ID 1220—LDAP over SSL" at technet.microsoft.com/en-us/library/dd941846(WS.10).aspx.

## Verifying LDAPS Connectivity

To verify that LDAPS is configured successfully on your DCs, you can use the LDP tool. LDP is installed by default on a Windows Server 2008 DC. On Windows Server 2008 member servers, Windows 7 machines, or Windows Vista machines, you must install Microsoft Remote Server Administration Tools (RSAT) to obtain access to LDP.

To open LDP, click Start and type ldp in the Search box. Click the Ldp Connection menu options, and then click Connect. In the Server field, enter the FQDN of the DC to which you want to connect. Ensure that the port is set to Port 636 (which is the default LDAPS port), that the Connectionless check box is cleared, and that the SSL check box is selected; then click OK. If LDAPS is configured properly, the LDP command output should display *Host supports SSL,* as in Figure 6.

Click the Connection menu option again, select Bind, and click OK. If LDAPS is configured properly, the LDP command output should display the username and domain name that you used for authenticating with LDP to AD.

## Closing Another AD Gate

Even though Microsoft doesn't provide a specific configuration interface, securing the LDAP traffic to your AD DCs using SSL/TLS technology is relatively easy. In this article, I explained how to enable LDAPS by installing a properly formatted certificate on your DCs. With LDAPS, you can lock down an important AD authentication and directory access gate. The two other main AD authentication protocols—Kerberos and NTLM—both leverage remote procedure calls (RPCs) for transport and have proper security and encryption mechanisms that are enabled by default.

InstantDoc ID 141170

**Jan De Clercq**
(jan.declercq@hp.com) is a member of HP's International Expertise Team and focuses on architecture for Microsoft-based IT infrastructures, identity management, cloud computing, and security. He's co-author of *Microsoft Windows Security Fundamentals* (Digital Press).

FEATURE

# Exchange Server 2010 Personal Archives

## Plan ahead before deploying archive mailboxes

by Tony Redmond

One of the most important new features in Microsoft Exchange Server 2010 is the introduction of Personal Archive mailboxes. Many third-party software vendors also provide archive solutions for Exchange. These solutions have worked for many years, and they often provide more developed features than Exchange 2010 offers in areas such as data ingestion and more sophisticated compliance functionality. In addition, third-party solutions typically don't focus only on Exchange but also offer the ability to archive data from other sources, such as Microsoft SharePoint, websites, and file shares. Finally, because they've stood the test of time, third-party solutions have real-life case studies and customer references about how to effectively manage the growth of information associated with Exchange over a sustained time period. But it's just too early to have the same degree of information about Exchange 2010 archiving.

Because archive mailboxes are built in to Exchange 2010 and the UI to access archives is available in the latest Microsoft clients, many companies—especially those that have never deployed archiving before—are attracted to the prospect of offloading information from primary mailboxes (the mailboxes used to send and receive messages) into an archive mailbox that's accessible online and can be managed through other Exchange 2010 compliance features, such as retention policies and multi-mailbox discovery searches. I can't think of any technology that can simply be introduced to a large group of users without some degree of planning—and deploying archive mailboxes without planning ahead is an especially bad idea.

### How Archives Work

When you enable a mailbox for an archive, you instruct Exchange to create a second mailbox that's marked for use as an archive. This mailbox can be held in the same mailbox database as the primary mailbox, or it can be assigned to a completely different database (in Exchange 2010 SP1 and later). It can also be situated in the cloud within a Microsoft Office 365 domain if you configure the hybrid connection required to link on-premises Exchange 2010 and Office 365. The separation of primary and archive mailboxes opens up design possibilities such as creating specific databases to store nothing but archive mailboxes (an archive database) or placing a set of archive databases on a dedicated mailbox server (an archive server). Both of these approaches have pros and cons, which are beyond the focus of this article.

An archive mailbox is simply another form of mailbox that's stored in a mailbox database. You can work with folders and items in the archive mailbox in exactly the same manner as you can with the primary mailbox. The primary and archive mailboxes for a user are linked through a globally unique identifier (GUID) maintained in the ArchiveMailboxGUID property of the user mailbox. GUIDs are 64-bit numbers that mean nothing to human beings. But in this case, they allow Exchange to locate a mailbox's archive no matter which database it's stored in. To discover which mailboxes in an organization have

archives, you can use the fact that mailboxes with archives have the ArchiveGUID property populated. Enter the following command, which looks for any mailbox where the link to an archive is not null:

```
Get-Mailbox -Filter {ArchiveGUID -ne
  $Null}
```

If you examine the archive-related properties that Exchange 2010 SP1 maintains for a mailbox, you'll see the database that holds the archive (ArchiveDatabase), its GUID, the name of the archive as displayed by clients, and the quotas that are used to control the point at which Exchange flags warnings about an approaching limit (ArchiveWarningQuota) and the point at which it's no longer possible to store more data in the archive (ArchiveQuota). The ArchiveDomain and ArchiveStatus properties are used only when an archive is stored on an Exchange server in the cloud running on Office 365.

## Clients and Archives

The introduction of archive mailboxes is a major upgrade for the Exchange Store, so it shouldn't come as a surprise to discover that not all clients are capable of revealing the presence of an archive or include the UI necessary to let a user interact with data held in the archive, including the UI to reveal retention policies and tags. In fact, full functionality is currently limited (at press time) to Microsoft Outlook 2010 or Outlook Web App (OWA).

In December 2010, Microsoft released an update for Outlook 2007 SP2 to let the Outlook 2007 client access archive mailboxes. This code works, and the archive mailbox is displayed like any other repository, such as a personal folder store (PST). However, Outlook 2007 doesn't include any of the UI features necessary to display data such as retention policies and tags—so some important functionality is invisible to users. In addition, Outlook 2007 doesn't perform searches automatically across primary and archive mailboxes in the same transparent manner as Outlook 2010 does. (For more information, see the Exchange team's blog post "Yes Virginia, there is Exchange 2010 archive support in Outlook 2007" at blogs.technet.com/b/exchange/archive/2010/12/20/3411710.aspx.)

Note that Autodiscover is the component that lets Outlook know about the presence of an archive mailbox for both Outlook 2010 and Outlook 2007. After you enable an archive mailbox, Autodiscover will detect its presence the next time Outlook starts up, and the archive will be automatically listed in the set of available repositories.

Neither Outlook for Mac 2011 nor any ActiveSync client supports access to archive mailboxes, probably because the underlying APIs (Exchange Web Services and ActiveSync) haven't yet incorporated the necessary API calls to open and manipulate archive mailboxes. You also can't work with archive mailboxes using a BlackBerry device because BlackBerry Enterprise Server doesn't support the necessary access.

All earlier versions of Outlook remain blissfully unaware of the archive but are able to access mailboxes on an Exchange 2010 server. POP3 and IMAP4 clients don't care about archives because these interfaces weren't designed to support a division of storage between primary and archive mailboxes.

Figure 1 shows how OWA presents an archive to a user. In this case, the user has opened a folder in the archive and is reading an item. The menu is revealed by a right click, and you can see that *Retention Policy* is shown to let the user select a retention tag to apply to the item. If an item were open in the primary mailbox, the user might also see *Archive Policy* listed in the menu to let the user apply tags to control when Exchange moves items into the archive mailbox.

## Exchange's Default Retention Policy

When you enable a mailbox to have an archive, Exchange automatically assigns a retention policy to the mailbox unless a retention policy is already assigned. The logic here is that the user who owns the mailbox can now move items into the archive. Exchange wants to make this process as easy as possible, so it makes sense that the user should be provided with some retention and archive tags to help manage the movement of items into the archive. Figure 1 provides a good example of how the tags in the policy are exposed by OWA.

Exchange's developers created a set of tags and gathered them in the *Default Archive and Retention Policy* that's created as part of the Exchange 2010 SP1 installation. This policy contains personal tags, which users can apply to items to keep them for a predetermined period of time before the Managed Folder Assistant moves the items into the Recoverable Items folder. The policy also contains archive tags, which dictate when items are moved into the archive mailbox.

Most of the tags in the *Default Archive and Retention Policy* are personal (shown

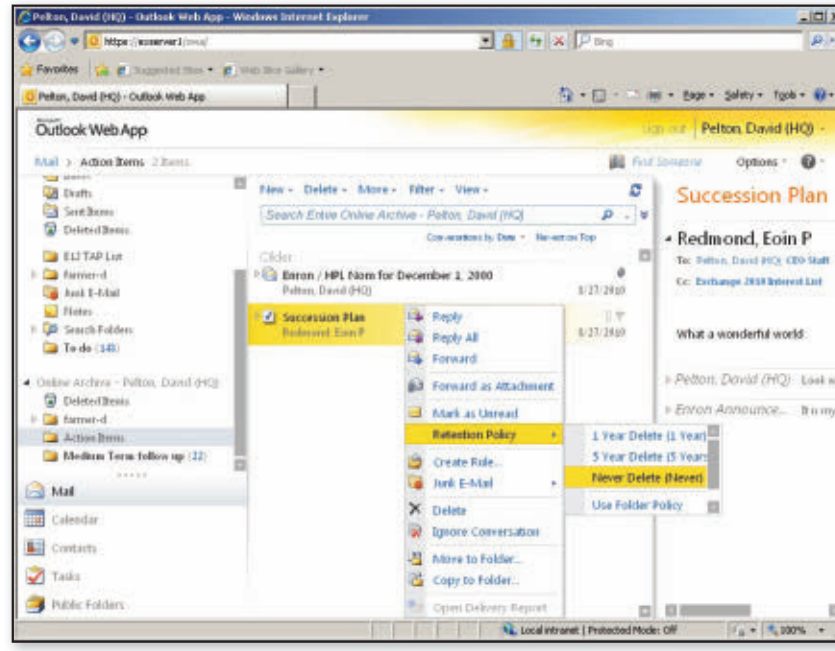


Figure 1: Accessing documents in an archive with Outlook Web App

as Personal Tag in Figure 2). In Exchange, this means that a tag must be explicitly applied by the user before the Managed Folder Assistant processes the action determined in the tag.

Tags contain a property called TriggerForRetention that determines the date used to calculate the age of an item. This date is typically measured from the date the item is first delivered to a mailbox; you can't alter or extend the date.

For example, if you select the *6 Month Delete* tag and assign it to an item, the Managed Folder Assistant moves the item to the Recoverable Items folder (the action DeleteAndAllowRecovery) after 6 months (the retention period), after first delivery to the mailbox (TriggerForRetention = WhenDelivered). The same is true for an archive tag such as *Personal 5 year move to archive.* In this case, the Managed Folder Assistant moves the item into the same folder in the archive (the action) after 5 years (the retention period). Retention periods are always stated in number of days, so 5 years is 1,825 days.

The highlighted tag shown in Figure 2 is named *Default 2 year move to archive,* which you can see applies to *All other folders in the mailbox.* This is a default policy tag (DPT), which means that the Managed Folder Assistant applies this tag to every item in the mailbox unless it comes under the control of another (more explicit) tag. We'll return to the effect that the DPT has on a mailbox shortly; it's sufficient for now to say that a DPT typically exerts a very powerful influence over a mailbox. Note that a retention policy can include two DPTs—one that influences when items are moved into the archive and the other that determines when items are moved into the Recoverable Items folder.

Exchange Server 2010 RTM displays a similar retention policy listed when you run the Get-RetentionPolicy cmdlet. This policy is called the Default Retention Policy. The difference between this policy and its SP1 equivalent is that the SP1 policy includes archive tags. The upgrade to SP1 leaves the old retention policy in place because it might have been applied to mailboxes. Removing the policy would invalidate the retention settings for user mailboxes, which would be a bad idea. However, this version of the default retention policy isn't much good because it contains a limited set of tags, so it would probably be best if you replaced it with a purpose-built retention policy designed to meet the business needs of the organization or the new default policy supplied with SP1.

Office 365 domains also use a default archive and retention policy called *Default MRM Policy* that includes a tag called *Default 2 year move to archive* that's applied to user mailboxes that have archives. Office 365 mailboxes are assigned archive mailboxes by default, so archiving is active immediately. Note that different Office 365 plans include different archive quotas. For example, the entry-level Office 365 Plan P1 for professionals and small businesses includes an archive quota of 25GB.

## The Role of the Managed Folder Assistant

The Managed Folder Assistant must process a mailbox before a user sees the effects of a retention policy, including the appearance of retention policies and tags in the Outlook 2010 and OWA UIs. It's important to communicate to users well ahead of time the fact that the Managed Folder Assistant affects mailbox contents because users will inevitably panic if they think items have been lost.

The Managed Folder Assistant can only do what it's told to do, and its instructions come from the retention tags that are placed on items. Users will understand that the Managed Folder Assistant processes an item according to a tag that the user explicitly places on it. For example, if you select an item and tag it to be retained for a year, you expect nothing to happen to that item during the year. Things get a little more dicey when the Managed Folder Assistant processes items according to the DPT, if one exists in a retention policy. Remember, a DPT dictates what happens to any item in a mailbox that isn't already under the control of another tag—so if the default tag for the retention policy applied to a mailbox dictates that items are deleted or archived after a set time period, that's what will happen.

The DPT in the *Default Archive and Retention Policy* specifies that items are moved into the archive mailbox after they are 2 years old (730 days). Most mailboxes hold items that are older than 2 years. It therefore follows that the Managed Folder Assistant has some work to do immediately after the policy is applied to a mailbox. In this case, the Managed Folder Assistant scans for any item that's over 2 years old and isn't stamped with another tag and moves it into the equivalent folder in the archive mailbox. This is extremely logical from a computer science perspective but incomprehensible for most users. Unless you tell users up front what will happen, they will perceive that items have disappeared from their mailbox. They won't think to look in the archive, where the Managed Folder Assistant has faithfully moved the missing data.

Of course, there are other consequences too. Anyone who uses Outlook configured in Cached Exchange Mode is accustomed to having replicas of all their server mailbox folders available in the OST and therefore accessible even when a network connection is unavailable to the Exchange server. However, items in archive folders aren't synchronized down to the OST and are therefore unavailable when the user is working offline. Again, this is logical because an archive is intended as a repository for information that's infrequently accessed, but that's probably not how the

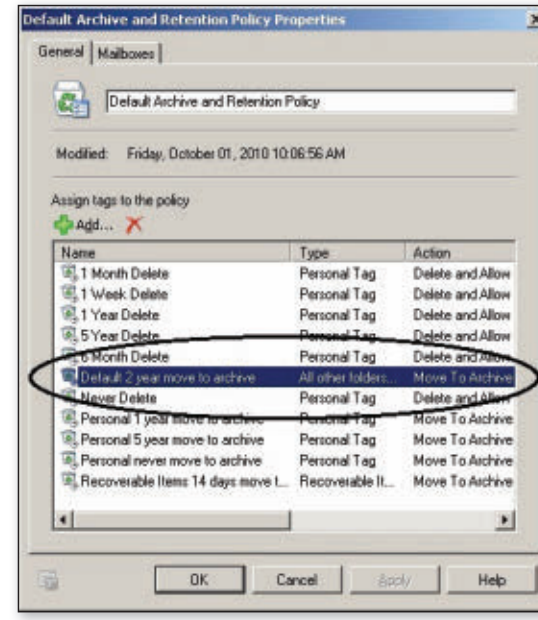


Figure 2: Tags in the *Default Archive and Retention Policy*

user who has just lost access to his data sees the situation.

It's also important to understand that the DPT exerts an ongoing influence over the mailbox and will be applied by Exchange to new items as they enter the mailbox, whether they arrive as new messages or are imported into the mailbox from a PST.

## Stressing Mailbox Servers

Applying retention policies to mailboxes can force the Managed Folder Assistant to do a lot of work as it processes items, especially if it has to move items from one database to another—as in the situation when the primary mailbox is in one database and the archive is in another. It's easy for an administrator to run a command to apply a retention policy to a group of mailboxes. For example, the following command creates a list of mailboxes in a specific database and then applies a retention policy to the mailboxes:

```
Get-Mailbox -Database 'DB1' |
  Set-Mailbox -RetentionPolicy
  'Management Retention Policy'
```

Running this command takes a matter of seconds, even if it has to process hundreds of mailboxes. By comparison, the Managed Folder Assistant has to crank into action and do a lot of work to process all those mailboxes. The Managed Folder Assistant in Exchange 2010 SP1 is more efficient at processing mailboxes because it schedules work automatically across the entire day rather than being constrained to a limited time window, as in Exchange 2010 RTM and Exchange Server 2007.

Formerly, the Managed Folder Assistant was given a fixed time period within which it could process mailboxes. Sometimes not all mailboxes could be processed within the set time period, leading to inconsistent results for users. Exchange 2010 SP1 includes the concept of work cycles, which is a way to assign workloads to Exchange components that must be processed within a certain time period. It's then up to Exchange to figure out how best to perform the necessary work within the allotted time period. In the case of the Managed Folder Assistant, its work cycle is 1 day, meaning that it's expected to process every mailbox on a server at least once daily. In effect, this means that the Managed Folder Assistant might be active throughout the day but will automatically throttle back its processing to respect the current workload of the server so that it doesn't interfere with the ability of the server to satisfy user demand. Interestingly, Microsoft uses a different work cycle for the Managed Folder Assistant on Office 365 servers to process mailboxes on a weekly basis.

Even with the more efficient work cycle, a lot of work must still be done to locate items that are now under the control of the policy and to apply the actions determined in the retention tags. As an example, let's assume that you apply the *Default Archive and Retention Policy* to 500 mailboxes and that each mailbox has 2,000 items that are over 2 years old. The Managed Folder Assistant must now process 100,000 items and move them into an archive.

Processing 100,000 items won't happen quickly, and it creates a reasonable demand for CPU and I/O on the mailbox server that hosts the user mailboxes, as well as on any other server that's associated with the activity, such as the servers that host the databases that contain the archive mailboxes. Given the world of multi-core high-end servers used to run Exchange 2010 and the automatic throttling used by the Managed Folder Assistant, it's difficult to be more precise about the workload generated on any particular server. However, it's fair to say that enabling archives can cause a spike in demand if the Managed Folder Assistant has many mailboxes to process for the first time. However, you might not notice the extra demand if this work occurs at night when the servers are otherwise not busy. If you run replicated database copies within a database availability group (DAG), additional resources are consumed to replicate and replay the transaction logs containing the transaction generated by Managed Folder Assistant activity on the servers that host the database copies.

## Processing Older Items

Another fact that users won't be aware of is that the Managed Folder Assistant applies the retention policy to all items in the mailbox. Again, this is totally logical because there's no point in assigning a retention policy to a mailbox unless Exchange ensures that its directives are respected. In most cases, because the DPT typically specifies a reasonably long retention period (2 years or more), the DPT doesn't have an immediate effect on items such as new messages that arrive in the user's Inbox. The retention countdown clock starts as soon as Exchange creates items in the mailbox, and the DPT will eventually move items into the archive or Recoverable Items folder—but only after their retention period expires. The situation becomes more interesting when you introduce older items into the equation.

Older items might be stored today in PSTs. I have items in PSTs that go back to the original Exchange Server 4.0 version shipped in 1996. Apart from laziness, I don't have a good reason why these items are still around. However, the point is that the typical user isn't good at cleaning out old items (which is why the need exists for the kind of automatic mailbox cleanup that you can implement with retention policies). Because users are human packrats, a lot of the items in their PSTs are probably not required and certainly should never occupy valuable space in an online database.

After you import data from a PST into a primary or archive mailbox, the Managed Folder Assistant examines the imported items the next time it processes the mailbox. It's probable that the Managed Folder Assistant will discover that a high percentage of items exceed the retention period specified in the DPT and will therefore apply whatever action is stated in the DPT. Thus, you can imagine a situation in which the following occurs:

1. The company decides that PSTs are difficult to manage and that a project will be created to import data from PSTs into Exchange 2010 mailboxes (primary or archive) so that all data in the company is accessible to multi-mailbox discovery searches. A similar approach might be taken if the company decides to migrate data from a third-party archive system used with a previous version of Exchange. Given that there are no migration features built in to Exchange 2010 to ingest data from third-party archives, it's likely that you'll have to take a two-phase approach to the migration and move data out of the third-party archive into PSTs and then import the data from the PSTs into

Exchange 2010. Third-party products are available from companies such as TransVault Software (www.transvault.com) and Sherpa Software (www.sherpasoftware.com) that can move items from different archives into Exchange. Because archiving products sometimes compress data, it's likely that some careful exercises in calculation will be required to estimate the data storage necessary to hold information as it passes from the third-party archive into Exchange 2010 SP1.

2. Administrators proceed to gather PSTs from users and import the data using the New-MailboxImportRequest cmdlet available in Exchange 2010 SP1. Note that this isn't an exercise that you should perform without careful planning because of the strain that the imports will exert on mailbox servers. It's also worth noting that you can import data directly into an archive mailbox with the New-MailboxImportRequest; no interim movement through the user's primary mailbox is necessary. Again, third-party software vendors have tools that can help you locate and ingest PST contents into Exchange 2010. Microsoft has promised to provide a PST ingestion tool (see the Exchange team's blog post "Coming Soon: PST Capture Tool" at blogs.technet.com/b/exchange/archive/2011/07/05/coming-soon-pst-capture-tool.aspx). However, that software hasn't yet appeared, even in beta—and even when it does, it's possible that the greater experience of the third-party software vendors in this space will mean that those vendors' tools will continue to be more functional and sophisticated. Useful abilities here include PST discovery on laptop disks, automatic removal of PSTs and blocking after their contents are imported, and policy-driven imports (e.g., import information only from the past 3 years and delete everything else that's found in the PST except in certain folders).

3. Users are happy to see all their PST data in their online mailboxes.

4. The Managed Folder Assistant runs to process user mailboxes and discovers all the items that have been imported from PSTs. The items are deemed to exceed the retention period specified in the DPT, and depending on where the item is stored (primary or archive mailbox) and the retention action determined by the DPT, the Managed Folder Assistant either moves the items into the Recoverable Items folder or into the archive mailbox.

5. Users now discover that some of their PST data is no longer where it was after the initial import and ask the Help desk what's going on.

Careful planning is necessary to understand exactly what happens to items after they're imported and how users might perceive the consequences. After you understand the flow, you can translate it into terms that users will understand and present what will happen in a positive and proactive manner.

## Successful Deployment

To successfully deploy archive mailboxes in Exchange 2010 SP1, you need to consider several factors. The following list is a good place to start.

1. Determine whether archive mailboxes satisfy the business needs of your organization. There's no point in deploying technology if it doesn't satisfy a business requirement. Remember that archives require enterprise CALs (eCALs), so factor this cost into the evaluation unless you already use eCALs for other purposes.

2. Determine the mailboxes that will be enabled with archives.

3. Determine whether the archive mailboxes will be in the same databases as their primary counterparts.

4. Determine what retention policies will be used within the organization and what policy will be assigned to different user groups. It's a good idea to assign the retention policies to mailboxes before you enable them for archives because this means that Exchange won't automatically assign the *Default Archive and Retention Policy* to the mailboxes.

5. Well before any retention policies are assigned, tell the people who own the mailboxes what the effect of the policies will be (e.g., when items will be moved to the Recoverable Items folder). Inform users about personal tags and how they can avoid the effect of the DPT by applying personal tags to individual items, complete folders, or conversations.

6. Make sure that the affected users have clients that reveal retention and archive tags. Outlook 2010 is best, OWA is acceptable, and Outlook 2007 is acceptable but has limitations.

7. A week before implementation day, remind users about the retention policies and their effect on mailboxes.

8. The day before implementation, send an email message to inform users that they might notice that some items have been moved into folders in the archive after the Managed Folder Assistant runs and processes their mailboxes. (Don't tell users about the Managed Folder Assistant; tell them what will happen in terms they will understand.) Explain how users can retrieve items from the archive and how easy it is to move items between the primary and archive mailboxes. Explain that items in the archive aren't accessible offline.

9. The day after implementation, make sure the Help desk is ready to handle calls from users who think they've lost items. Provide Help desk staff with a 1-2-3 cheat sheet for dealing with users that explains how to find items in the archive and how to move them back into the mailbox and stamp them with an appropriate personal tag to prevent the Managed Folder Assistant from moving them back into the archive again.

10. Take a well-earned rest and prepare for the next group of users.

## Be Prepared

Archive mailboxes are a great new feature in Exchange 2010—but like so many new features, the mere fact that technology is available doesn't mean you can simply deploy it with your brain in neutral. Some thought and careful planning will ensure that both you and your users survive the introduction of archive mailboxes with just a few scars.

InstantDoc ID 140655

**Tony Redmond** (12knocksinna@gmail.com) is a contributing editor for *Windows IT Pro*, and author of *Microsoft Exchange Server 2010 Inside Out* (Microsoft Press). He blogs about Exchange and related topics at www.windowsitpro.com/go/ExchangeUnwashed.

# Use PowerShell to Report on Scheduled Tasks

## An alternative to the Schtasks utility

by Bill Stewart

The Windows NT family of OSs has had a built-in program scheduler since its inception. It has grown from the command-line based At scheduler available in Windows NT to the more powerful Task Scheduler service that debuted in Windows 2000. In Windows Server 2008 and Windows Vista, Microsoft overhauled the Task Scheduler service and provided even more functionality.

Windows 2000 didn't provide a script object or command-line interface to the Task Scheduler service, with the exception of the difficult-to-use command-line utility Jt.exe, which was provided in the Microsoft Windows NT Server Resource Kit and not the OS. However, this situation improved starting in Windows XP, which provided the Schtasks utility.

I've seen various requests in newsgroups and online forums asking for a way to create a report of scheduled tasks on one or more computers. I'll describe the OS's built-in way of doing this (namely, Schtasks.exe with the /query parameter), why its output format is difficult to use for reporting purposes, and how I solved this problem using a PowerShell script.

### The Problem with Schtasks

As I mentioned previously, XP and later includes the Schtasks utility, which is a command-line interface to the Task Scheduler. The Schtasks command's /query parameter outputs a list of scheduled tasks on a computer. For example, the command

```
schtasks /query /s server1 /fo CSV
```

outputs the scheduled tasks on the computer named server1 in comma-separated value (CSV) format, which is suitable for importing into a spreadsheet or database. The Schtasks /query command works fine on XP and Windows 2003, but if you use it on later versions, you'll run into problems. Vista, Server 2008, and later versions support task folders, and unfortunately the Schtasks command outputs a separate CSV header row for each task folder, even if a task folder doesn't contain any tasks. Figure 1 shows a sample of CSV data imported from the *Schtasks /query /fo CSV* command, with the repeated CSV header rows highlighted. It isn't a big problem to delete the extra rows from the output for one computer, but this doesn't scale well when you need to report on scheduled tasks for many computers.

The other Schtasks output formats have problems of their own. The List format (/fo List) provides a newline-separated list, but this format is difficult to parse. The Table format (/fo Table)

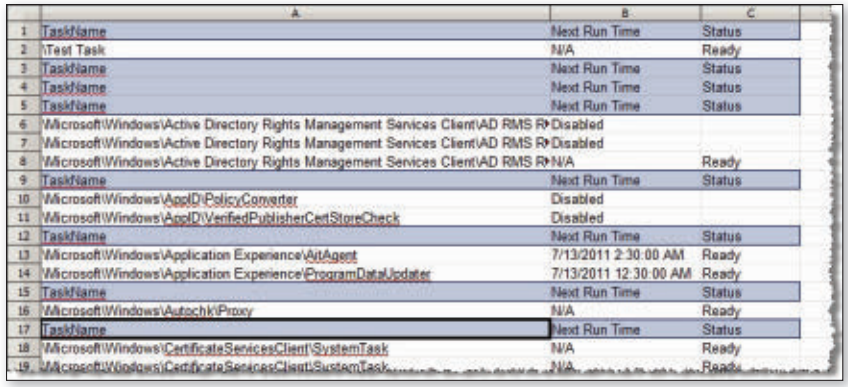

| | A | B | C |
|---|---|---|---|
| 1 | TaskName | Next Run Time | Status |
| 2 | \Test Task | N/A | Ready |
| 3 | TaskName | Next Run Time | Status |
| 4 | TaskName | Next Run Time | Status |
| 5 | TaskName | Next Run Time | Status |
| 6 | Microsoft\Windows\Active Directory Rights Management Services Client\AD RMS R | Disabled | |
| 7 | Microsoft\Windows\Active Directory Rights Management Services Client\AD RMS R | Disabled | |
| 8 | Microsoft\Windows\Active Directory Rights Management Services Client\AD RMS R | N/A | Ready |
| 9 | TaskName | Next Run Time | Status |
| 10 | Microsoft\Windows\AppID\PolicyConverter | Disabled | |
| 11 | Microsoft\Windows\AppID\VerifiedPublisherCertStoreCheck | Disabled | |
| 12 | TaskName | Next Run Time | Status |
| 13 | Microsoft\Windows\Application Experience\AitAgent | 7/13/2011 2:30:00 AM | Ready |
| 14 | Microsoft\Windows\Application Experience\ProgramDataUpdater | 7/13/2011 12:30:00 AM | Ready |
| 15 | TaskName | Next Run Time | Status |
| 16 | Microsoft\Windows\Autochk\Proxy | N/A | Ready |
| 17 | TaskName | Next Run Time | Status |
| 18 | Microsoft\Windows\CertificateServicesClient\SystemTask | N/A | Ready |
| 19 | Microsoft\Windows\CertificateServicesClient\SystemTask | N/A | Ready |

Figure 1: Repeated CSV header rows in the Schtasks command's output

supports a /nh (no headers) option, but there are blank lines in the output and the output is separated by task folders, making it difficult to parse as well. The XML format (/XML) requires writing XML-parsing code to generate a usable report.

Rather than wrestle with parsing the Schtasks command's output, I decided to look for a better way. Fortunately, you can use the Task Scheduler scripting objects (msdn.microsoft.com/en-us/library/aa383607.aspx) to get scheduled task information. I decided to write a PowerShell script, Get-ScheduledTask.ps1, that uses these objects to output scheduled tasks on one or more computers.

## Introducing Get-ScheduledTask.ps1

Get-ScheduledTask.ps1 requires Vista, Server 2008, or later because the Task Service object isn't available on earlier OS versions. You must also run the script from an elevated PowerShell session. (To do this, right-click the PowerShell icon and choose *Run as administrator.*) The script's syntax is as follows:

```
Get-ScheduledTask
[[-TaskName] <String[]>]
[[-ComputerName] <String[]>]
[-Subfolders]
[-ConnectionCredential <PSCredential>]
```

The -TaskName parameter specifies the name of one or more scheduled tasks. Wildcards (* and ?) are allowed. You can also specify a list of task names separated by commas or a variable containing an array. If you omit this parameter, the default is "*" (i.e., output all tasks). You can omit the -TaskName parameter name if its argument is first on the command line.

The -ComputerName parameter specifies the name of one or more computers. Wildcards aren't allowed with this parameter, but you can specify a text file that contains a list of computer names (one name per line). This parameter also supports pipeline input. If you don't specify a computer name, the current computer is the default. You can omit the -ComputerName parameter name if its argument is second on the command line.

The -Subfolders parameter specifies whether the script supports task subfolders. Without this parameter, the script works only with tasks in the root tasks folder ("\"). In the Task Scheduler GUI, the root tasks folder is the topmost folder in the hierarchy. If a remote computer doesn't support task folders, this parameter is ignored.

The -ConnectionCredential parameter requires a PSCredential object (created with the Get-Credential cmdlet) that contains the credentials you want to use to connect to the TaskService object on the specified computers. Please note that this is a potentially insecure operation. The script must get a plaintext copy of the PSCredential object's password because the TaskService object's Connect method doesn't support encrypted credentials.

Get-ScheduledTask.ps1 outputs custom objects (PSObjects) for each task. Table 1 lists the default object properties output by the script. If the default set of properties provides too much information, you can use the Select-Object cmdlet to select only the properties you want. If a property isn't supported (e.g., a property that doesn't exist on an older OS), it will be empty. Also, if the ActionType is not Execute, the Action property will be empty.

Table 2 shows sample commands to run Get-ScheduledTask.ps1. Note that although

Table 1: Default Properties of the Output Object in Get-ScheduledTask.ps1

| Property | Data Type | Description |
|---|---|---|
| ComputerName | String | Specifies the computer hosting the scheduled task. |
| ServiceVersion | String | Provides the internal version of the Scheduled Task service (e.g., "1.3"). |
| TaskName | String | Specifies the task's name (including the path, if supported). |
| Enabled | Boolean | Is True if the task is enabled; otherwise is False. |
| ActionNumber | Int | Is an arbitrary number the script assigns to differentiate each output object. The first action in the task is 1, the second action is 2, and so on. |
| ActionType | String | Specifies the type of action: Execute (runs a command), COMhandler (fires a handler), Email (sends an email), or ShowMessage (shows a message box). OS versions earlier than Vista/Server 2008 support only Execute. |
| Action | String | Contains the command line if ActionType is Execute. |
| LastRunTime | DateTime | Specifies the date and time the task was last run. |
| LastResult | String | Provides the hexadecimal result code (0x0 is success) for the task the last time it ran. |
| NextRunTime | DateTime | Indicates the date and time the task is scheduled to run next. |
| State | String | Specifies the task's operational state: Unknown, Disabled, Queued, Ready, or Running. |
| Author | String | Indicates the user account used to create the task. |
| Created | DateTime | Specifies the date and time the task was created. |
| RunAs | String | Indicates the user account used to run the task. |
| Elevated | Boolean | Is True if the task is configured to run with highest privileges or False if not. |

Table 2: Sample Commands to Run Get-ScheduledTask.ps1

| Command* | Description |
|---|---|
| Get-ScheduledTask *time* | Outputs scheduled tasks in the root tasks folder on the current computer that contain the word *time*. The -TaskName parameter name isn't required because it's the first parameter on the command line. |
| Get-ScheduledTask -ComputerName server1,server2 | Outputs the scheduled tasks on two computers (server1 and server2). |
| Get-Content Computers.txt \| Get-ScheduledTask -Subfolders | Outputs all scheduled tasks, including those in subfolders, for each computer listed in Computers.txt. |
| Get-ScheduledTask \Microsoft* computer1 -Subfolders | Outputs tasks matching the pattern \Microsoft* on computer1. The -TaskName and -ComputerName parameter names aren't required because they're the first and second parameters on the command line. The -Subfolders parameter is required in this case because the task name pattern contains a task folder name (\Microsoft). |
| Get-ScheduledTask -Subfolders \| Export-CSV Tasks.csv -Encoding ASCII -NoTypeInformation | Exports all scheduled tasks on the current computer to the Tasks.csv file. |
| Get-ScheduledTask -ComputerName server1 -Subfolders -ConnectionCredential (Get-Credential) | Outputs all scheduled tasks on server1. The parentheses around the Get-Credential cmdlet cause PowerShell to execute it as an expression. The expression's result (i.e., a PSCredential object) is used for the -ConnectionCredential parameter. |
| Get-ScheduledTask \| Select-Object TaskName,LastRunTime | Outputs tasks in the root tasks folder on the current computer and sends this output to the Select-Object cmdlet, which outputs only the TaskName and LastRunTime properties. |

* Although the commands wrap here, you'd enter them on one line in the PowerShell console.

Listing 1: Code That Creates the TaskService Object

```
try {
  $TaskService = new-object -comobject "Schedule.Service"
}
catch [System.Management.Automation.PSArgumentException] {
  throw $_
}
```

the commands wrap in the table, you'd enter them on one line in the PowerShell console.

## How Does It Work?

The Get-ScheduledTask.ps1 script is designed to operate like a cmdlet because it can use pipeline input in place of the -ComputerName parameter. To support this ability, the script uses the *begin* and *process* script blocks. In the *begin* script block, the script defines some script-wide (global) variables, then attempts to create an instance of the TaskService object, as shown in Listing 1. If the script can't create the object, it throws an error and ends. After creating the TaskService object, the script defines all of the supporting functions. The workhorse function of the script is the get-scheduledtask2 function, which I'll describe in a moment.

The *process* script block is where the script passes each computer name to the get-scheduledtask2 function. The $PIPELINEINPUT variable, defined at the top of the *begin* script block, enables the *process* script block to determine whether it should use the -ComputerName parameter or retrieve its input from the pipeline.

## The get-scheduledtask2 Function

The get-scheduledtask2 function is the workhorse function of the script, as I mentioned previously. The *process* script block executes this function for each computer name. If the -ConnectionCredential parameter is used, the function extracts the domain name, username, and plaintext copy of the password in the PSCredential object to pass to the TaskService object's Connect method. If the Connect method fails, the function outputs a warning message using the Write-Warning cmdlet and exits from the function.

After this, the get-scheduledtask2 function checks the TaskService object's HighestVersion property to determine the version of the Task Scheduler service. The function uses the service's version number to determine whether there's support for task folders and to determine the value for the output object's Elevated property (see Table 1).

The function then uses the TaskService object's GetFolder method to get the root task folder, after which it uses the get-task function to retrieve a list of all tasks on the computer. Next, the get-scheduledtask2 function compares the task's name with the name in the -TaskName parameter. If there isn't a match, the function continues to the next task. If there is a match, the function sets the $actionCount variable to zero and iterates through the task's Actions collection.

For each action in the task, the get-scheduledtask2 function uses the New-Object cmdlet to create an empty custom object (a PSObject) and adds the task's properties to PSObject. After adding all the properties, it outputs the object.

## Know What's Scheduled on Your Computers

You can use Microsoft's Schtasks utility to report on scheduled tasks, but its output formats are limited, difficult to parse, and don't scale well. Get-ScheduledTask.ps1 overcomes these limitations and makes it easy to report on scheduled tasks for as many computers as you need. You can download this script by going to www.windowsitpro.com, entering 140978 in the Search box, and clicking the 140978.zip hotlink.

InstantDoc ID 140978

**Bill Stewart**
(bstewart@iname.com) is a scripting guru who works in the IT infrastructure group at Emcore in Albuquerque, New Mexico. He has written numerous articles about Windows scripting, is a moderator for Microsoft's Scripting Guys forum, and offers free tools on his website at westmesatech.com.

# Paul Thurrott...


## . . . he's not in Microsoft's pocket, but now he can be in yours.

The independent voice for IT enthusiasts

Paul Thurrott delivers news, tips, commentaries, and reviews on Microsoft technology – from gaming to mobile to servers to software, and coverage of Microsoft competitors in between. Get daily updates without reaching farther than your pocket.

Download your
Paul Thurrott: PocketTech app today

**windowsitpro.com/mobile-apps**

Available for iPhone | Windows Phone 7 | Android

# On-Premises vs. Hosted Email Archiving

## Evaluate the pros and cons of various solutions

by Paul Robichaux

Some topics are polarizing by their very nature. For example, if you ask a random person about a reality TV show such as *So You Think You Can Dance*, you'll likely get one of two responses: The person will either love the show or hate it. Email archiving fosters the same kind of polarized reaction. People who work for organizations that have archiving and compliance requirements are intensely interested in the topic, whereas people who work for organizations that don't have such requirements typically pay little or no attention to archiving.

However, this polarization is slowly changing as more organizations realize the benefits of archiving corporate email messages. These advantages include ease of compliance with regulatory requirements and improved productivity for end users who benefit from having all their mail available at once.

Microsoft has been fairly aggressive in the archiving space, introducing its own on-premises archiving solution in the form of Exchange Server 2010's Personal Archive feature. Exchange 2010 SP1 allows a mailbox and its associated archive to be in different mailbox databases, which means your mailboxes can be hosted by Microsoft Office 365 and your archives can be hosted on your own servers, or vice versa. Several archiving vendors also sell their own hosted archiving systems. Evaluating the pros and cons of each type of archiving system will help you understand the benefits and drawbacks of each, so that you can determine how a particular solution might meet your own needs.

### Do You Need Archiving?

The first question to ask is whether you actually need archiving in the first place. There are three possible answers to this question: Yes, No, and Maybe.

Some companies are required by law, regulation, or other considerations to maintain archives of their email messages. Examples abound; for example, financial services companies in most countries are required to archive at least some of their electronic communications, and many choose to go beyond the basic requirements to reduce their liability.

A small number of companies absolutely do not want to use archiving because they don't want a corporate-sponsored repository of valuable data that becomes subject to legal discovery or regulatory inquiry. To the extent that they're required to use archiving for regulatory or compliance reasons, they aggressively limit what they archive, how long it's kept, and who's allowed to set or change archiving policies.

Most companies fall between these extremes. They don't have a defined legal or regulatory requirement to archive their email messages. Instead, they might archive messages to reduce the cost of messaging services, to prepare for potential discovery or compliance requirements in the future, or to improve employee efficiency by giving users a robust archiving mechanism. Organizations in this

group can often categorize the demand for archiving based on the stakeholders who are asking for it:

- IT stakeholders are primarily concerned with cost reduction, data management, and service provision. If you're asking questions such as "Will archiving help extend the useful life of my servers?" and "Can archiving reduce the cost of my primary Exchange storage?" then you fall into this category.
- Business stakeholders are primarily concerned with efficiency and productivity. For them, archiving is a way to help users be more productive by allowing fast recovery of required data.
- Legal stakeholders are typically tasked with keeping the organization and its employees out of jail and out of the headlines. Their archiving requirements involve the ability to quickly and accurately respond to discovery requests, dispose of unneeded data, and set policies that ensure that the requirements they face are met consistently.

Because of the overlap between these requirements, you might think that selecting an archiving solution would be straightforward. In practice, what typically happens is that Exchange administrators are asked to choose a solution based on their own understanding of the requirements—which often leads to buying an inappropriate solution. If archiving is important to your business, you must think of it as a long-term strategy, not a check box or a purchase order.

## On-Premises Pros and Cons

The market for on-premises archiving solutions is mature because archiving solutions have been around nearly as long as email server products have. As email infiltrated the financial services, pharmaceutical, and government sectors, customers demanded robust archiving solutions—so major archiving vendors have had years to build solutions. In fact, the market has evolved to the point where email archiving itself is only a small part of most products' feature sets. For example, it's common for archiving products to provide tools for ingesting unstructured data such as file shares and Microsoft SharePoint libraries, integration with case and litigation management systems, and other bells and whistles that go beyond basic archival and e-discovery features.

The major advantage of on-premises archiving is that you're in complete control. You have both complete authority and complete responsibility. You get to choose what's archived, where it's stored, who has access to it, and so on. However, if problems develop with the archive, there's nowhere to point the finger of blame. For example, if you're required to perform a discovery search as part of a court case and you can't produce all the necessary records, you probably won't be able to blame the vendor.

You should also keep in mind that a high degree of control also requires a high degree of operational maturity and experience. Even the best-designed, easiest-to-use systems require some administrator time—and a poorly designed or complex system requires that much more. If you don't have the time or in-house knowledge required to manage a full-blown on-premises archiving system, then a hosted offering might be a better choice.

If archiving is important to your business, you must think of it as a long-term strategy.

Another advantage of on-premises archiving is that on-premises systems tend to have much greater functionality than hosted systems. On-premises systems can ingest and manage more types of data, given the fact that most organizations don't want to make all their data externally accessible to hosted archiving tools. In addition, on-premises systems provide tighter integration with a wider variety of back-end systems—and many of them provide customization capabilities as well. In general, if you need to archive SharePoint, file server, or other types of data besides email, you probably need an on-premises solution.

The pros and cons of on-premises solutions are more mixed when it comes to cost. You typically must purchase the entire archiving system up front, which means that for most organizations, archiving is initially funded as a capital expenditure rather than from operating funds. The need to purchase all the services you require means that it's somewhat more difficult to deploy pilot programs with on-premises systems, because you must buy all the major components in order to get even a single mailbox archived—which can be a major barrier to adoption unless you're absolutely certain which archiving product you want to deploy.

## Hosted Pros and Cons

Hosted services of all kinds share a few common attributes. One is that they tend to offer pay-as-you-go pricing. This makes them attractive to customers who want to be able to predict exact costs for the services they use. Keep in mind, of course, that the hosting provider can (within the limits of whatever contract you negotiate) change the price for archiving services. Many hosting providers price their services according to the amount of archive data you store—which certainly seems reasonable, although it puts a premium on your ability to estimate how much storage you'll use over the term of your hosted service contract. However, this disadvantage might be small compared with the flexibility of being able to lease or subscribe to the services you need for the term in which you need them.

Another aspect of hosted archiving services that makes them attractive is their ease of deployment. Typically, hosted archives let you feed them your mailbox data over time. Web-based archiving systems let you deploy archive search-and-discovery facilities to users who need them without installing or configuring desktop client software, which is another significant benefit.

Hosted services put the burden of management, maintenance, monitoring, and security on your hosted service provider. If you want "set-it-and-forget-it" archiving capability, hosted solutions can give it to you. It's a good idea to carefully review the archiving provider's service level agreement (SLA), and of course you should

## Learning Path

### *Windows IT Pro* Resources

"Exchange 2010 Architecture: Microsoft's Ankur Kothari Talks About Personal Archives," InstantDoc ID 129821

"Exchange 2010 Archive Mailbox Sizing and Scaling," InstantDoc ID 140894

"Planning for Exchange Server 2010 Personal Archives," InstantDoc ID 140655

"Real-World Exchange 2010 Migration: Implementing the New Stuff," InstantDoc ID 136611

### Microsoft Resources

"Email Archiving and Retention," www.microsoft.com/exchange/en-us/email-archiving-and-retention.aspx

"Exchange Online - Hosted Email for Business," www.microsoft.com/exchange/en-us/exchange-online-hosted-email.aspx

"Introduction to Personal Archive," office.microsoft.com/en-us/outlook-help/introduction-to-personal-archive-HA101830421.aspx

"Understanding Personal Archives," technet.microsoft.com/en-us/library/dd979795.aspx

thoroughly investigate a hosted archive provider's customer references before signing an agreement. Run—don't walk—away from any vendor that makes it difficult for you to do either of these things.

## Exchange 2010's Archiving

In your consideration of on-premises archiving systems, be sure to include Exchange 2010's built-in archiving features. Although these features might not be a perfect fit for every organization, they're priced into Exchange 2010—which might let you hit some of your archiving requirements with minimal additional expense.

Microsoft has made a lot of noise around its Software Plus Services (S+S) strategy, but Exchange 2010 and Office 365 are a great example of how this strategy can pay off. You can host your mailboxes on your own Exchange servers and store archives on Exchange Online, or vice versa. You have a great deal of flexibility regarding how you manage and operate your mailboxes and where you store your data. Other archiving vendors can't yet match the tight integration between Exchange Online and on-premises Exchange. It's reasonable to expect a continuing movement to more tightly integrate these services (along with the other components of Office 365).

Keep in mind, though, that Exchange's archiving doesn't do everything that larger, more complex archiving products do. Its strengths lie in cost, integration (such as the fact that discovery searches are controlled by Exchange's Role Based Access Control—RBAC—permissions feature), and ease of deployment. If you have complex discovery requirements, you might find that a third-party solution, whether on-premises or hosted, is a better fit for your needs.

## Considerations

As you evaluate whether a hosted or on-premises solution is best for your environment, you should ask numerous questions. The following five questions are especially important:

> Exchange Server's archiving strengths include its cost, integration, and ease of deployment.

1. Why am I archiving? Understanding the business reasons that underlie your archiving requirements is critical to choosing the right combination of products and services. If there are specific laws or regulations that you must meet, you need to know what they are and what they require. Your organization's legal and business stakeholders are important sources to help answer this question.

2. How predictable is my deployment? Do you have good information about how much data you need to archive and how it's likely to grow in the future? Without these figures, you might lean toward a fixed-cost on-premises solution, but with better information you can make a more informed decision about the subscription or lease cost of a hosted solution.

3. Can I handle success? If your archiving system gets more use than expected or requires more administrative time than you plan for, will you have the resources to handle it?

4. How much do I trust my archiving vendor? The archiving marketplace has been undergoing consolidation for the past 2 or 3 years, and that trend appears set to continue. It's important that you're comfortable with your archiving vendor's history of product support, track record for support of new Exchange releases, and road map for supporting on-premises, hosted, and hybrid solutions.

5. What's my roadmap? In other words, what kind of data growth do you expect for all the data types you have to archive? Are there business changes, such as mergers, acquisitions, or entries into new business areas ahead that might change your archiving needs? What about Exchange upgrades and migrations? This question is a great opportunity for you to assess any other future factors that you think might influence your choice of product, hosting mode, or deployment schedule.

If you're not certain of the answers to any of these questions, consider how you could pilot archiving solutions to help provide more fodder for making an informed decision.

## The Future of Archiving

The brisk competition between hosted and on-premises services in general is sure to continue. To decide between a hosted and on-premises solution for email archiving, you must understand your archiving requirements and how each type of solution might meet those requirements. As Microsoft begins work on the successor to Exchange 2010, and as archiving vendors continue to merge and consolidate, expect to see significant changes in this space within the next 2 or 3 years.

InstantDoc ID 140496

**Paul Robichaux**
(probichaux@windowsitpro.com) is a senior contributing editor for *Windows IT Pro*, a content author at Acuitus, and a Microsoft Exchange Server MVP and MCSE. Paul is the author of *Exchange Server Cookbook* (O'Reilly and Associates) and blogs at www.robichaux.net/blog.

FEATURE

# 10 Reasons *Not* to Brand SharePoint

Think that branding is a must? You might be surprised!

by Michael T. Smith

It seems that one of the first things people want to do with a new Microsoft SharePoint installation is to brand it. Branding public-facing SharePoint sites is considered practically mandatory. Branding internal corporate portals to reinforce the company image might also make sense. But the most common use of SharePoint within an organization is for departmental sites, team-collaboration sites, and document-management sites. Should you brand these internal sites?

There are two kinds of SharePoint branding for internal sites. One preserves the full SharePoint UI and feature set. This simple branding modifies graphics, colors, and font types. It uses features that are built in to SharePoint to let site owners update site navigation and Web Parts. This branding might involve changes to Cascading Style Sheets (CSS) or edits to the SharePoint master pages, but it leaves the UI completely predictable to the average SharePoint user and can be easily supported without outside help.

Anything more complex falls into the second category of branding. This type of branding often involves an outside branding consultant and hours upon hours of planning, design, and implementation to match the external company website or an older, custom internal site. This type of branding changes how SharePoint and its UI work. Before you decide to brand internal sites by using this second category of customizations, ask yourself the following 10 questions. (If you still insist on branding your SharePoint installation after reading this article, see the web-exclusive sidebar "If You Must Brand SharePoint," www.windowsitpro.com, InstantDoc ID 141138.)

## #1: Would you pay to brand Windows Explorer or Microsoft Excel?

Have you branded your word processor or your email client? Of course not! These are tools. They should have a consistent and predictable UI, such as an obvious start button. After learning how to use a tool one time, you should be able to figure how to use the same kind of tool the next time. SharePoint is also a tool, especially when used for team collaboration and document management. Branding sites that are used for those purposes—especially when users might access more than one site—should be treated as such.

## #2: Do you want to increase your per-user costs?

The per-user cost of a SharePoint installation is fairly reasonable. That is, until you start spending $10,000 to $30,000 per department—or even per site—to pay for a graphics design firm or branding consultant to customize your internal sites. The real-world branding costs can easily be in the hundreds of dollars per user and provide only a cosmetic benefit.

## #3: How fast do you want users to get to work?

Customizing UIs takes time and often delays a new SharePoint installation. When branding has been approved, teams are put together to get the sites branded. These teams must interview consultants, review designs, wait for delivery, and test the result before the sites can be deployed to users. And then, if each site looks different, with a different and unpredictable UI, users will be wasting time figuring out how to navigate the site and how to find content.

## #4: How much do you want to spend on training?

Out of the box, SharePoint has a wealth of available training and support resources, including instructor-led classes, books, online videos, and endless web resources. All these resources are affordable (or even free) but are useful only for uncustomized sites. Custom UIs require custom training; without it, users are less productive.

## #5: How much do you want to spend on support?

If each site is different, will your support groups be able to help your site users? Will your Help desk be able to answer questions such as, "In the HR site, I click on a green duck to get to the employee manuals, but I just went to the IT site to find software manuals, and there's no green duck. There are just two trucks, a race car, and a go-cart. Which should I click?" (If you think the duck-and-cars example is ridiculous, I'm not just being silly. I've seen many branded SharePoint sites that can be described only as "unique.")

This brings up a related issue: Graphic designers aren't always good SharePoint designers. Graphic designers tend to think of SharePoint as just another custom website and often break or remove the most basic features, such as Quick Launch or the ability to add or change a Web Part. After the consultant, designer, or brander has finished with the site, who will pay for fixing such issues, or even updating the site later?

## #6: How much time do you want to waste?

Often, the site owner is the one doing the branding. SharePoint Designer is free, easy to download, and talked about everywhere on the web. And it's so easy to use that site owners often become self-taught site web designers, spending much of their time playing with SharePoint Designer. This problem isn't new. Remember the early spreadsheet days, when managers switched from managing teams to spending all day playing with spreadsheets? In the age of SharePoint, we have managers and team leaders spending too much time as web designers. Most of these site owners have no design training and no governance.

## #7: Do you know who's in charge?

When every department is doing its own thing with SharePoint, is any department doing the right thing with corporate assets? Are site owners following corporate standards for site content and content governance? If you lose control of SharePoint and the content that's stored there, you might never get it back (short of starting over from scratch). And when the legal or R&D departments ask, "Can you find X?" or "Can you tell me who did Y?" are your SharePoint sites organized and structured enough to actually perform an audit?

## #8: How difficult will sites be to audit?

If each department and team feels free to create custom UIs as a means of branding, then they also might feel free to store their content any way they like. If they have their own branding, then they will surely have their own custom content types, list types, and metadata. How will a researcher or auditor find anything in such a system? Imagine being an auditor who must visit a hundred sites, each with a different UI, to find a document about a customer or a product. This Wild West approach is expensive and difficult to maintain.

## #9: Are there better places to invest?

How much sense does it make to try to reduce costs by licensing SharePoint Foundation or SharePoint Server Standard Edition, only to spend a lot of money on custom (and cosmetic) branding, and then more money on custom training and lost productivity because of the branding? For the same price, you can stay with out-of-the-box SharePoint and spend the extra money on SharePoint Enterprise Edition, Microsoft FAST Search Server, and some powerful business intelligence (BI) tools. You might even have enough to invest in faster hardware or more user training. If you're interested in doing things the right way, then invest in a governance plan and an ongoing governance team.

## #10: Do you really want to do this all over again?

Your branding costs don't end with the current installation of SharePoint. Sooner or later, along comes the next generation of SharePoint with a whole shopping cart full of new features that you want and need. Branded sites almost never upgrade cleanly. Over the past few years, I've seen how the migration from SharePoint 2003 to SharePoint 2007—and more recently from SharePoint 2007 to SharePoint 2010—has worked for branded sites. Typically, it hasn't been a good experience and has required paying to rebrand all the sites to work in the new version. Are you willing to bet on the effort and cost of moving your branded sites to the next version of SharePoint?

## The Bottom Line

Before you make the decision to brand internal sites, make sure you have a real business need to do so. Talk to other companies that use SharePoint, and find out what it's really costing them to brand sites, including the ongoing support costs. Will new hires be able to figure out the custom UIs and site designs? Can you afford the up-front costs, ongoing support costs, end-user training costs, and eventual upgrade costs of branding? And what about legal and business accessibility requirements (e.g., support for screen readers, high-contrast text, nonmouse navigation, Web Content Accessibility Guidelines—WCAG—2.0). How might branding affect these requirements? In a nutshell, do you really need to brand?

InstantDoc ID 141137

**Michael T. Smith**
(mikesmith@microsmithinc.net) is an instructor at MAX Technical Training, a SharePoint MVP for 2010 and 2011, and a Microsoft Certified Trainer (MCT). He trains, consults, and writes about SharePoint administration, development, and governance.

- Infragistics
- Twisted Pair
- Opscode
- Acronis

## Infragistics Debuts NetAdvantage for SharePoint

Infragistics released NetAdvantage for SharePoint 2011, a set of SharePoint Web Parts that enables business professionals to build code-free dashboards to quickly visualize Key Performance Indicators (KPIs) and actionable metrics. In NetAdvantage for SharePoint, Infragistics introduces the Map Web Part, enabling SharePoint users to map anything and everything in an eye-catching and easy-to-read fashion. The Map Web Part is ideal for professionals who want to display data at multiple levels such as a street, city, state, country, or any geographic coordinate for a visual experience that viewers can quickly interpret and take action from. With the Data Chart and Gauge Web Parts, users can easily extract meaning from the numbers and present clear renditions of important key performance indicators and high-end business intelligence. See the NetAdvantage for SharePoint page at www.infragistics.com for a complete list of new features.

## Twisted Pair Improves Mobile Communication for SharePoint Users

Twisted Pair Solutions announced the availability of WAVE Communicator for Microsoft SharePoint, an application for integrating communication capabilities into SharePoint 2007 and 2010. By adding secure voice capability along with status, presence, and location information to a SharePoint user's experience, WAVE Communicator improves information flow between office-based and mobile workers, increasing cross-functional collaboration and decision making. Built on the WAVE communication interoperability platform, WAVE Communicator for SharePoint includes a Web Part that provides IT departments with the ability to add mobile worker communications and collaboration directly into a new or existing SharePoint deployment. WAVE Communicator also makes it possible for SharePoint users to reference presence information for mobile workers equipped with WAVE Mobile Communicators and GPS-supported two-way radios, enabling team members to locate field resources immediately and get the right resources assigned to the task at hand. Contact Twisted Pair Solutions at www.twistpair.com.

## RPost Debuts Integrated Security, Legal, and Document Services for Microsoft Outlook

RPost released a major upgrade to its RPost desktop software extension for Microsoft Outlook. RPost for Outlook 2000-2010 is an integrated messaging platform that provides for legal electronic messaging, business productivity, unique messaging collaboration, email encryption security, and document signing services—all offered through a light plug-in that adds a special Send button to the Outlook toolbar for on-demand functionality. As part of this upgrade, RPost also offers a developers' version to its underlying Outlook software that programmatically launches RPost's Outlook UI and functionality. This permits developers to add RPost functions without any Outlook programming complexity. The upgraded extension also offers support for Outlook 2010 32-bit and 64-bit, extensive configurability for IT staff pre-installation or

PRODUCT SPOTLIGHT

## Opscode Delivers Cloud Infrastructure Automation to Windows Environments

Opscode announced that its Chef software and commercially supported Hosted Chef and Private Chef will now provide infrastructure automation in Windows environments. Opscode's release enables broad deployment and automation of key components of Windows infrastructure, including PowerShell, Internet Information Services (IIS), SQL Server, and Windows services.

Opscode Chef is an open-source systems integration framework built for automating the cloud. It allows IT teams to easily deploy thousands of servers and scale applications throughout an entire infrastructure. Through a combination of configuration management and service-oriented architectures, Chef, Hosted Chef, and Private Chef make it easy to create an elegant, fully automated infrastructure. Customers can use Opscode Chef to configure raw machines as web servers, and then to manage the web application deployment, automating all the core components of big web shops on Windows.

With support for Windows infrastructure, Chef software—known as cookbooks—can now provide the setup, automation, and maintenance of Windows-based servers and applications, while still allowing companies that have deployed PowerShell to leverage their current investment. Opscode Chef allows administrators to include PowerShell into cookbook recipes, which Chef then propagates automatically, enabling consistent and repeated use and distribution in large-scale environments.

Opscode Chef also fully supports PowerShell and SQL Server, and it allows for IIS installation, deployment, and configuration. For more information about Opscode Chef, contact Opscode at www.opscode.com.

by end users after installation, and security features for IT staff to lock in feature settings. Contact RPost at www.rpost.com.

## LANState Pro 6.0 Offers Visual Network Monitoring with SNMP Trap Support

10-Strike Software announced LANState Pro 6.0, a Windows-based network monitoring and management utility with a web UI. The program allows network administrators and engineers to make a graphic map of a network and see the current state of each host in real time. LANState Pro monitors servers, switches, databases, processes, folders, files, disks, installed software, and other objects and devices. The program alerts the systems administrator when the hosts and services go down and displays the monitoring results on the graphic map. LANState Pro 6.0 includes more than 20 new features and bug fixes to help users monitor and control network computers and servers more efficiently. The most important new features include SNMP trap support, Windows Event Log monitoring, and response time charts in the Web UI. Find a complete overview of the new features and improvements at www.10-strike.com/lanstate/history.shtml.

## Acronis Updates Snap Deploy

Acronis released Acronis Snap Deploy 4, the latest version of its automated solution for deployment of data and systems to multiple PCs, virtual machines (VMs), and servers. Acronis Snap Deploy 4 helps IT administrators and lab managers copy exact images of data and systems (including the OS, applications, and files) and move it from one device to another. Organizations operate a mix of PCs, servers, and OSs, which drives up the cost of management, including the cost of deployment. Acronis Snap Deploy 4 enables IT departments to standardize and control their infrastructure by deploying a standard master image, and reduce costs and time by automating deployment. Key features include per-machine customization, hands-free mass deployment with centralized management and automation, and hardware-independent images. Pricing starts at $25 per workstation and $121 per server. Contact Acronis at www.acronis.com.

## RSA Offers Secure User Access to the Microsoft Cloud

RSA announced that its RSA SecurID multi-factor authentication solution can be integrated with Active Directory Federation Services (AD FS) 2.0. This integration lets organizations provide secure user access to Office 365 cloud-based web applications such as Microsoft Exchange Server technology. Secure, federated access control enables an organization to extend the protection for remote users to applications deployed outside an organization's network. RSA SecurID and AD FS 2.0 work together to help provide secure access to cloud-based or remote services with identities provisioned and administered by enterprise IT. RSA SecurID technology solves the "weak link" issue of poorly chosen passwords by enforcing strong, multi-factor authentication. RSA SecurID technology also helps simplify the user experience, resulting in greater efficiency of identity and access management and support, which can ultimately result in lower costs. Contact RSA at www.rsa.com.

# Paul's Picks

www.winsupersite.com

**SUMMARIES** of in-depth product reviews on Paul Thurrott's SuperSite for Windows

### Samsung Focus S

**PROS:** Windows Phone 7.5; AMOLED screen; excellent camera; tethering support

**CONS:** A bit large for some users; not really 4G

**RATING:** ◆◆◆◆◇

**RECOMMENDATION:** Samsung's flagship device for Windows Phone 7.5 is a winner, with a large but thin and light form factor, a gorgeous 4.3-inch Super AMOLED Plus display, 16GB of storage, and a speedy 1.4GHz processor. It sports both front- and back-facing cameras, and the primary camera is unexpectedly excellent, with 8 megapixels of resolution and true-to-life colors. It brings the best-in-market mobile OS, Microsoft's Windows Phone 7.5, which includes Internet sharing (aka tethering) via AT&T's wireless network. However, despite its airy lightness, it might be a bit big for some customers, and AT&T's claimed 4G speeds are a far cry from the real thing.

**CONTACT:** Samsung • www.samsung.com/us/mobile/cell-phones/SGH-I937OKAATT

**DISCUSSION:** www.winsupersite.com/article/windowsphone75/windows-phone-75-samsung-focus-141245

### Samsung Focus Flash

**PROS:** Windows Phone 7.5; small form factor; tethering support; inexpensive

**CONS:** So-so camera and storage; not really 4G

**RATING:** ◆◆◆◇◇

**RECOMMENDATION:** The Focus S is an obvious follow-up to last year's Focus, but the Focus Flash is different. First, it costs just $50 with a new 2-year wireless contract. (The Focus S is a more standard $200.) Second, it comes in a small, light, and fun form factor, with tapered sides, a gorgeous but small 3.7-inch AMOLED screen, and, in a curious first, a mechanical Start button surrounded by capacitive Back and Search buttons. It stumbles a bit in the specs, though the form factor will attract those already on the fence about large-sized smartphones. Aside from its powerful 1.4GHz processor, it offers a lackluster 8GB of storage, a middling 5 megapixel camera, and AT&T's not-really-4G wireless network support. But it's so cute, none of that might matter.

**CONTACT:** Samsung • www.samsung.com/focusflash

**DISCUSSION:** www.winsupersite.com/article/windowsphone75/windows-phone-75-samsung-focus-flash-141268

InstantDoc ID 141274

REVIEW

# Colligo Contributor Pro 4.3

SharePoint, in one form or another, has been one of Microsoft's most successful products in recent years, and has become accessible to organizations of all sizes. Out of the box, SharePoint offers users a web interface, which works well but involves something of a learning curve for most users, and basic integration with Microsoft Outlook, for offline synchronization of SharePoint libraries. The limitations of these built-in components can pose a problem with SharePoint user adoption, especially if Microsoft Exchange Server public folders and file servers serve the same purpose.

Microsoft Office Pro Plus 2010 introduced SharePoint Workspace, an application that's based on the software that Microsoft acquired from Groove Networks. SharePoint Workspace provides a desktop-client front end for SharePoint, along with some of the functionality that was provided in Microsoft Groove. Although SharePoint Workspace is a step in the right direction, it's available only to those who license Office Pro Plus and has some significant shortcomings.

If you're looking for a more advanced level of functionality, Colligo's Contributor Pro provides full SharePoint integration with Outlook and Windows Explorer, through the Contributor Add-In for Outlook and Contributor File Manager. The included desktop client—Colligo Contributor Client—completes the suite.

## Contributor Add-In for Outlook

Colligo Contributor Add-In for Outlook synchronizes not only SharePoint document libraries but also custom and standard lists such as events, contacts, and agendas. Unlike when using the standard Outlook interface, this add-in allows you to add files to SharePoint. However, probably the most significant feature is the ability to drag email from Outlook into a SharePoint library and have Contributor copy all the important metadata, such as sender, subject, and date sent. This capability is vital when moving email out of Exchange to another storage medium: Metadata provides the basis for most user searches and is crucial in ensuring that the data remains discoverable. If you use Outlook's default SharePoint integration side-by-side with the Contributor Add-In for Outlook, you'll soon discover just how limited the standard Outlook functionality is.

Being able to drop email into SharePoint is all very well, but if doing so is to be a manual operation, it's unlikely to serve as a long-term storage solution that frees up valuable space in your Exchange database. With the Contributor Add-In installed, Outlook rules can be used to automate the filing of email to SharePoint. Often, it makes sense to store the important information that's contained within email in a location that's easily accessible to all members of a project team and that can be managed according to corporate policy. If necessary, you can drag an email attachment alone to SharePoint, leaving the original message in Outlook, with a link to the document.

The add-in's Send & File button replaces the standard Send button in Outlook, so users can have all sent items automatically filed in SharePoint. Colligo Contributor also supports the copying of entire Outlook folders to SharePoint, with full metadata capture. Custom metadata and content types can also be configured. And when creating a new email in Outlook, you can choose to upload documents to SharePoint and send the SharePoint links rather than attachments—a useful feature for reducing storage costs. You can edit metadata directly in Outlook by right-clicking an item in a SharePoint library, and you can use views to organize data in the same way as in the SharePoint web interface.

## Contributor File Manager and Contributor Client

Colligo Contributor File Manager provides drag-and-drop support for SharePoint in Windows Explorer. Contributor Client is a standalone application that includes the functionality of the other programs and allows you to filter synchronization with SharePoint. This capability can save time when only specific documents need to be accessed from a SharePoint library. As with Contributor Add-In for Outlook, File Manager allows users to check in or check out documents and edit properties directly from Windows Explorer. Documents and folders can be moved between SharePoint sites, and email messages can be dragged from Outlook to Contributor Client.

## Integration Ideal

Most SharePoint-integration features are available across all three products in the suite. The question is which interface users prefer to work with; it's unlikely that all three applications would be used on a regular basis. The Contributor Configuration Editor allows system administrators to preconfigure synchronization for SharePoint sites, synchronization frequency, and other important settings.

Although the individual applications worked as advertised, I did have a couple problems when switching between them. Colligo is aware of these issues and told me that they will be fixed in an upcoming service pack and in the next release of File Manager. Contributor is an expensive solution, but if you're using SharePoint extensively, the price can be offset against reduced training costs and productivity improvements, no doubt making this suite an invaluable addition to the SharePoint experience.

InstantDoc ID 141139

### Colligo Contributor Pro 4.3

**PROS:** Complete solution integrates SharePoint with applications that users are familiar with

**CONS:** Price

**RATING:** ◆◆◆◆◇

**PRICE:** Starts at $179; volume discounts apply after 10 seats

**RECOMMENDATION:** If your budget can stretch to it, Colligo Contributor Pro provides a valuable SharePoint add-on that can help reduce training costs and improve productivity.

**CONTACT:** Colligo Networks • 866-685-7962 • www.colligo.com

Russell Smith | rms45@rsitc.com

# FastTrack Scripting Host

If I have a systems management problem, I like to use a bit of scripting to provide a solution. But scripts aren't always the best solution; they can be complex, and unless the rest of the IT staff also has good coding skills, any required modification can pose a problem. For that reason, I always recommend that IT shops look at Group Policy and Group Policy Preferences before investigating custom scripted solutions.

Group Policy Preferences offer advanced functionality that can replace a legacy logon script, but there will be occasion to apply a custom solution. FastTrack Scripting Host (FSH) is a replacement for native VBScript and Windows PowerShell. Although VBScript and the Windows Script Host (WSH) are not especially difficult to learn, a lot of development time and testing by someone who has the skills to work with and debug the scripting language is required when tasks become more complex. As for PowerShell, it is finding its place in newer versions of Windows but is harder to learn and closer to real programming than VBScript is.

Though much of what FSH does can be achieved by using VBScript or PowerShell, there's no doubt that FSH is faster in initial coding and testing. The product also provides useful workarounds for problems that are specific to the kind of scripting tasks that systems administrators commonly undertake. The development environment uses a drag-and-drop interface to build code, with helpful descriptions of syntax and function at the bottom of the main window. The product guides you through forming each line of the script. This "one task, one line of code" paradigm makes FSH easy to learn, but that's not to say that it doesn't require any understanding. I was able to compile some basic scripts without reading up on the syntax, but that might not be the case for those with little or no scripting experience.

It's easy to include attractive GUI elements in FSH scripts, beyond the basic dialog boxes and text input that VBScript offers. You also have the option to add icons, menus, buttons, and other elements. To run a FSH script from your company's Netlogon share, all you need to do is include FSH.exe and the .fsh script in the share. No client-side component is required. However, if you use the free FastTrack Logon component, the necessary files will be automatically copied from the Netlogon share to the user's application directory and run locally, for more efficient execution.

FSH supports what it refers to as offline scripts. A small executable—SmartDock—runs in the background, detects IP address changes, and then triggers a script to run. This capability is especially useful for notebook users who connect to different networks or via a VPN, should the user environment need to be customized for specific network connections.

One of the biggest issues with traditional scripting methods, such as VBScript, is that if you need to call an executable or run a routine by using credentials other than those of the logged-on user, you'll either need to prompt the user to enter an alternate username and password—which is hardly ideal—or need to embed the credentials in the script, a method that is insecure. FSH allows systems administrators to distribute scripts, either in native .fsh format or as .exe files, with embedded credentials that use strong encryption. Even notebook users without administrative privileges can be sent a quickly created and tested script that runs with admin privileges to install a piece of software or carry out an essential piece of maintenance.

Although other products can achieve the same result, FSH doesn't rely on any components, other than the .exe that the end user runs. FSH also contains a series of commands to simplify dealing with User Account Control (UAC) prompts under Protected Administrator accounts (i.e., administrator accounts that have UAC enabled).

Another powerful feature is the ability to distribute scripts as Windows Installer (.msi) or standard .exe files. Although interesting in its own right, this ability results in a much more intriguing capability: to repackage—or wrap—software installers as .msi files without using snapshots. If you've ever created an .msi installer file from before-and-after system snapshots, for use with a software distribution system such as Group Policy or System Center Configuration Manager (SCCM), then you know how hit-and-miss the results can be. Repackaging generally requires a lot of testing and a certain amount of experience. FSH scripts can be used to repackage programs without bypassing the installer logic, which is key to reliable deployment. Any elevation of privilege can be handled from within the installer itself. Windows Installer packages that you create by using FSH support the Windows Installer Major Upgrade flag, allowing automatic redeployment if a new package is detected.

Many of the tasks for which FSH might be deployed can be accomplished in other ways, and in many cases you should stick to standard technologies such as Group Policy Preferences, if they meet your needs. Nevertheless, Group Policy Preferences have their limitations. If you have tricky problems to solve, such as elevating privilege securely from inside a script or repackaging legacy applications, then this product might be the answer.

InstantDoc ID 140986

## FastTrack Scripting Host

**PROS:** Decreases development and testing time compared with native scripting languages; addresses solutions to common systems management problems

**CONS:** Requires a non-standard scripting language

**RATING:** ◆◆◆◆◆

**PRICE:** $499 for 50 licenses; contact vendor for pricing for more than 50 licenses

**RECOMMENDATION:** If Group Policy Preferences don't provide enough flexibility for tailoring user environments, or if you must regularly provision and modify custom scripts, then FastTrack Scripting Host provides a good alternative to VBScript and Windows PowerShell.

**CONTACT:** FastTrack Software • 888-446-7898 • www.fasttrackscript.com

Russell Smith | rms45@rsitc.com

REVIEW

# Messageware OWA Desktop

When you're planning a migration to Microsoft Exchange Server 2010, don't forget to plan for the client through which your users will access email. Messageware OWA Desktop gives users desktop access to Microsoft Outlook Web App (OWA) in Exchange 2010, allowing you to forgo the desktop Outlook client altogether and avoid the cost of getting your users off older Outlook or non-Outlook client versions.

OWA Desktop runs on Windows 7, Windows Vista, and Windows XP, as well as on Windows Server 2008 and Windows Server 2003. Installation is simple and wizard-driven: The whole process took less than 1 minute on my Windows 7 laptop. The first time you run OWA Desktop, the Create New Account window opens so that you can connect to an Exchange email account.

You have the option of manual or automatic configuration. For automatic configuration, you give your account a nickname, and then enter your Exchange email address, username, and password. OWA Desktop retrieves your server settings, and you're set to go. If automatic configuration can't find your account, manual configuration lets you input your OWA server URL and configure other advanced options.

When OWA Desktop is running, its icon appears in the Windows system tray. You can hover the mouse over the icon to see how many unread messages you have. Click the icon to open OWA Desktop Commander, which gives you complete control over your email environment—everything from reading and composing messages, to adjusting your settings, to opening a full OWA screen.

Part of the beauty of this product is the quick, easy access to email functions that might take several steps in OWA or Outlook. By using Commander, you can select Compose Mail, Compose Appointment, or Compose Task to open a new email message, appointment, or task. Each function provides formatting controls, access to your Global Address Book (GAL) and Contacts, and other features that you'd get if you were logged on to OWA.

Also in Commander, you can select View Unread to open a window with all your unread messages. This window has controls for reading, replying, forwarding, deleting, and marking a message read or unread. The View Reminders function shows you all current reminders, including calendar appointments and any reminders that you've placed on email messages. From this window, you can open the original item or dismiss or snooze the reminder.

You can use Commander to open a full copy of your Inbox, Calendar, or Tasks list. You can even launch a full OWA screen. The launch is quick, and for all Commander functions, OWA Desktop handles the sign-in in the background. The OWA screen gives you the complete OWA experience, including access to your account settings through the Exchange Control Panel (ECP).

The Commander pop-up has options for OWA Desktop account settings. You can run multiple email accounts, which each appear as an icon on the task bar, for individual control. You can set options for notifications and reminders, and you can access OWA Desktop Help, a well-written, interactive HTML file.

A unique feature of OWA Desktop is the Import and Export command, which has two primary functions. First, the command lets you import holidays to your calendar. You can include multiple nationalities on your calendar—handy if you work with international clients. Second, the command lets you import or export Contacts. A wizard-driven interface walks you through the process of either operation.

OWA Desktop works with Microsoft Office 365 as well as traditional onsite environments. For enterprise deployments, you can work with Messageware support to get a custom Deployment Package that lets you control which features are available to end users. For instance, users with an Office 365 Kiosk Workers (K1 or K2) plan might need to access email but have no need for Tasks or Calendar. Or you might want to lock down account, import, or export functionality, even for onsite users.

While I was testing OWA Desktop, I traveled to a conference in another time zone. I reset the time zone on my laptop, expecting my calendar to pick up the change, as Outlook does. I realized that I had a problem when my smartphone (which automatically adjusts to the time difference) and my OWA desktop notifications were out of sync. After a bit of research, I discovered that the problem was with OWA, not OWA Desktop. OWA, as a server-based system, doesn't get its time from my local computer, and I had to go into my OWA settings in ECP to reset the time zone. That action allowed OWA Desktop to pick up the change, and everything worked smoothly thereafter.

Similarly, any criticism I have of OWA Desktop ultimately goes back to the limitations of OWA itself. You can't use an image in your email signature, for instance, or add words to the user dictionary. Outlook's advanced formatting features are missing as well. It would be nice to see OWA Desktop develop the ability to do some of the things that OWA can't.

Nonetheless, OWA Desktop provides an easy-to-use, lightweight interface for OWA, and though it won't satisfy power users, it's a slick, cost-effective way to access OWA. If you're using Office 365 or moving to Exchange 2010 but don't want the expense of upgrading Outlook, OWA Desktop is a great option to consider.

InstantDoc ID 141265

## Messageware OWA Desktop

**PROS:** Full OWA experience on the desktop without repeated sign-ins; quick access to most-used features; easy to install, set up, and use

**CONS:** Limited by OWA itself

**RATING:** ◆◆◆◆◇

**PRICE:** $1.25 per user per month

**RECOMMENDATION:** If you're using Office 365 or you're moving to Exchange 2010 but don't want the expense of upgrading Outlook, OWA Desktop is a slick, cost-effective option to consider.

**CONTACT:** Messageware • 905-812-0638 • www.messageware.com

B. K. Winstead | bwinstead@windowsitpro.com

# Enterprise Random Password Manager

At some stage in your career, you've probably been handed the password to another administrator's account—written on a piece of scratch paper—for use in some unusual set of circumstances. It would be nice to think that, as soon as that password is used for its intended purpose, the person responsible for managing it would update it so that it's secret once more. In the trenches, though, what probably happens is this: Six months later, when a similar situation arises, the person says, "Mate, it's the same password I gave you last time!"

Enterprise Random Password Manager (ERPM) is designed to deal with the issue of privileged password management in cross-platform enterprise environments. ERPM allows you to manage privileged Windows accounts, Linux or UNIX privileged accounts, service accounts, and application-specific accounts.

## Sophisticated Process

ERPM works in a way that's more sophisticated than just remembering the passwords for specific accounts. When an IT pro needs to perform a task that requires the use of a privileged account, he or she logs on to the ERPM web console and requests a password for that account.

Depending on how you configure EPRM, the request might be approved automatically, or the IT pro might need to wait for approval. Either way, when the request is approved, ERPM will issue the IT pro a complex temporary password for the account. This password can be displayed on the screen, sent through email, or transmitted though a text message.

ERPM ensures that the password has been synchronized on the related system before issuing it to the IT pro. Unlike typical administrator passwords, this password is valid for a limited time only; it then expires and the password is reset. Administrators also have the option of checking in a password, at which point the password will be reset ahead of schedule.

You can associate ERPM with a service desk application such as Microsoft System Center Configuration Manager (SCCM), thereby ensuring that an appropriate approvals framework is in place before passwords for sensitive accounts are dispensed.

## Simplified Management

The benefit of ERPM is that it simplifies the management of privileged accounts. Organizations can more easily keep track of who has access to privileged account credentials because those credentials are checked out for a specific amount of time only. Even if an IT pro changes the temporary password, ERPM will still reset it when the checkout period expires. Rather than having access to privileged accounts on an ongoing basis, administrators have access only when they need it to perform their designated job role.

After being configured in an environment, ERPM uses a continuous discovery process to find and secure new privileged accounts. For example, if you've deployed Microsoft SQL Server and add new accounts and databases, ERPM is updated with these credentials as they are created.

Many organizations use simple passwords for inter-application communication, as a way of simplifying the process. The problem with this approach is that simple passwords are more likely to be compromised by attackers. Because application and service passwords are managed centrally through ERPM, they can be substantially more sophisticated.

ERPM can also determine application interdependencies and update credentials accordingly. This capability solves one of an IT pro's biggest headaches: rotating service account and application passwords. When this rotation is performed manually, one or more account instances are inevitably forgotten and the service or application stops working because of authentication issues.

## Advanced Reporting

The other substantial advantage of ERPM is that because it uses a SQL Server back end, you can generate sophisticated auditing and compliance reports that show which passwords have been checked out of the system, by whom, and for what purpose. ERPM supports password management for as many as 120,000 systems with as many as 360,000 accounts per system—for a total support of as many as 3 million accounts.

## Special Caveat

My only concern about ERPM is, to mix a metaphor, about placing all the keys to the castle in the same basket. As ERPM can change any password in your organization, the administrator who controls the product indirectly controls everything. Special care must be taken when setting up ERPM, to ensure that it's secure. If incorrectly configured, the product could represent a large and tempting weak link in the organization's security infrastructure.

## Change for the Better

ERPM provides a logical framework for the management of privileged account credentials. Although the change to using temporary administrator passwords (rather than long-term, non-expiring passwords) will come as a bit of a culture shock to many IT pros, it can—if properly implemented—make privileged account management more auditable and secure.

InstantDoc ID 140979

### Enterprise Random Password Manager

**PROS:** Provides an easy-to-use method of automating the management and security of sensitive privileged accounts in heterogeneous environments

**CONS:** Requires that IT pros to take the time to adapt to its methodology

**RATING:** ◆◆◆◆◆

**PRICE:** $25,000 for a 500-workstation/device license

**RECOMMENDATION:** ERPM provides an effective solution for minimizing the distribution of passwords for sensitive accounts.

**CONTACT:** Lieberman Software • 800-829-6263 • www.liebsoft.com

Orin Thomas | orin@windowsitpro.com

# WE KNOW YOUR INDUSTRY.

AGRICULTURE

AUTOMOTIVE & TRUCKING

BROADCAST ENGINEERING

AVIATION

COMMERCIAL REAL ESTATE

DESIGN ENGINEERING

ELECTRONICS

ELECTRICAL SYSTEMS, ENERGY & CONSTRUCTION

FOOD

HEALTHY LIFESTYLE

IT & DEVELOPER

MANUFACTURING & SUPPLY CHAIN

MARKETING & MEETINGS

MECHANICAL SYSTEMS

PUBLIC INFRASTRUCTURE

RESTAURANTS

WEALTH MANAGEMENT

# WE KNOW YOUR CUSTOMERS.

WHAT THEY'RE READING. WHAT THEY'RE SEARCHING FOR. WHAT THEY'RE DOWNLOADING.

# WE KNOW HOW TO CREATE AND DELIVER CONTENT THAT PERFORMS.

WHITE PAPERS

WEBINARS

CONFERENCES

SEARCH MARKETING

SOCIAL MEDIA

LEAD NURTURING

VIDEOS

WEBSITES

RESEARCH

Penton Marketing Services™
WE KNOW YOUR CUSTOMERS
powered by eyetraffic™

Windows IT Pro

**CONTACT:**
Chrissy Ferraro
970-204-0952
**PentonMarketingServices.com**

# Self-Service Password Reset Managers

## These products can help end users help themselves

by Nate McAlmond

For large and growing companies, the task of assisting end users can become a tremendous burden on the IT department. By some estimates, the cost of password resets can be as much as $70 per incident (including loss of productivity) and make up around 30 percent of Help desk calls. Even higher costs can be expected in industries that are subject to additional regulation, such as in the financial and healthcare arenas.

### Product Similarities

All the products that I compared installed on a single Windows Server 2008 system in about 30 minutes or less. My installations each included an administrative console for configuring the software, an end-user website that users could use to reset forgotten passwords, and a Help desk website that Help desk workers could use to assist end users with password resets. Each product also checked passwords as they were entered and enforced a set of password requirements. The password requirements of all five products were similar; the only major exceptions were the dictionary options in Specops Password Policy and Quest Password Manager, which allows you to configure these products to prevent the use of specific words in passwords.

By some estimates, the cost of password resets can be as much as $70 per incident and make up 30 percent of Help desk calls.

The products' security features also had several similarities. Each product used a password-protected enrollment process, during which the end user completes a series of questions: You can require some questions or configure the products to present end users with a list of questions to choose from. All the reviewed products had rules to force end users to answer these questions in a useful and secure way. These rules included such options as

- requiring unique answers to all questions
- requiring answers to questions to be case sensitive
- setting the number of allowed custom questions
- setting the total number of questions
- requiring end users to set up password reset questions and to complete the enrollment process when it presents itself at logon
- setting a lockout threshold for incorrect answers to password reset questions (similar to lockout thresholds for password input during logon)
- setting a minimum custom-question length
- requiring all answers to be more than five characters
- restricting answers from including words that are in the question

Only ManageEngine's ADSelfService Plus did *not* use Microsoft IIS. Each product also included a client application that added a logon assistance button to the Windows logon screen. By clicking this button, end users are brought to a self-service password-management portal, without needing to log on to the computer. Without the client application, end users can still access the password reset website for enrollment into the system or to reset passwords. However, users who need resets will probably need to use a coworker's computer or a kiosk computer that allows web access without logging on first.

Another nice feature of the products is that they are licensed per user rather than per server. This feature allows you to set up a second server for fault tolerance.

### Specops Password Policy and Specops Password Reset

Specops Software's Specops Password Policy (which Figure 1 shows) and Specops Password Reset, two products that I tested together, install on Windows Server 2008, using the straightforward, checklist-like installation wizards that I've seen in other Specops products. As you move through each point of the installation, the wizards either take care of the requirement for you or tell you what needs to be done before proceeding. Within about 15 minutes each, installation was completed. These products are the only ones that install a self-signed certificate during installation to help secure all your web traffic. (The vendor recommends replacing the self-signed certificate with one from a public source after you move past the trial phase of your implementation. Otherwise, internal and external users will receive warning messages as they use the web-based self-service portal.)

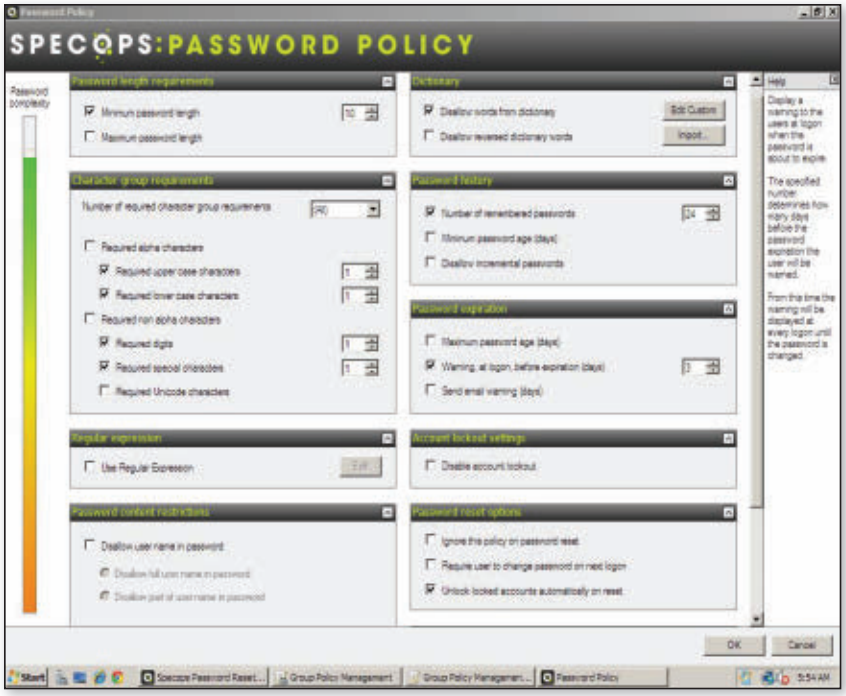


Figure 1: Specops Password Policy

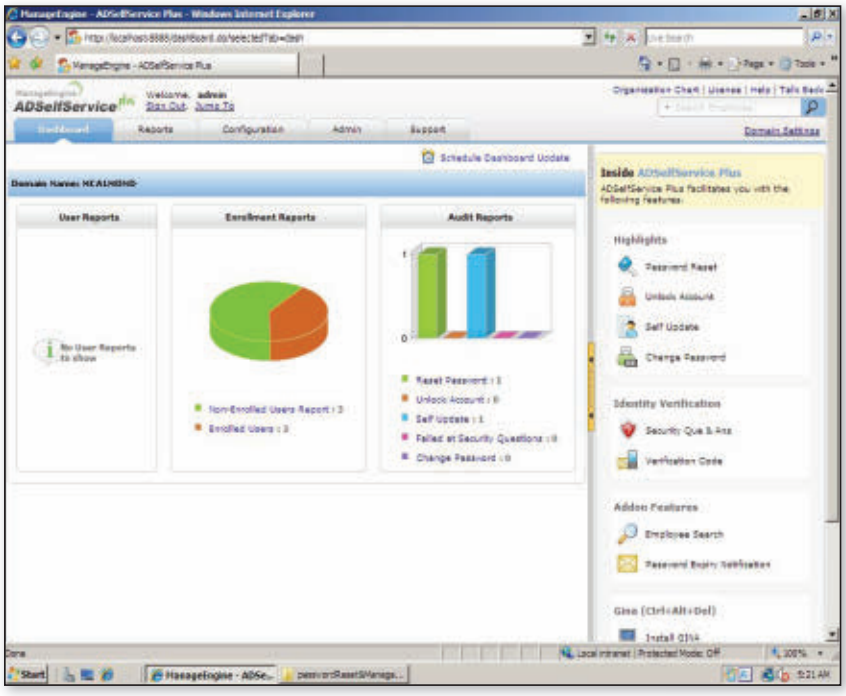


Figure 2: ManageEngine's ADSelfService Plus

You will need to install the included Specops Password Policy Sentinel on all DCs to which your users can connect, to ensure that Specops Password Policy is enforced during password changes throughout your organization.

When you set up the password-reset and password policies in Active Directory (AD), you need to set your domain password-policy requirements low and apply rigorous password requirements to individual organizational units (OUs). The password policies that you create with the Specops products need to be more restrictive to be compatible with the Default Domain Policy that is applied across the domain.

## Specops Password Policy and Specops Password Reset

**PROS:** Easy to deploy; integrates with AD to allow for multiple password policies; gives visual indication of password-policy strength; shows the user a dynamic view of the rules, indicating which requirements the password meets

**CONS:** No significant disadvantages

**RATING:** ◆◆◆◆◆

**PRICE:** $12,960 for 1,000 users ($6,480 for Specops Password Policy and $6,480 for Specops Password Reset)

**RECOMMENDATION:** These two Specops products provide an extremely well-put-together solution with a strong focus on security and completely integrate with AD, to the point of using native Group Policy management tools to configure the products.

**CONTACT:** Specops Software • 877-773-2677 • www.specopssoft.com

## ManageEngine's ADSelfService Plus

ManageEngine's ADSelfService Plus (which Figure 2 shows) has a full set of features for the price, plus an employee directory that users can use to search for other users' contact information or to update their own. The product installs with just a few clicks and uses its own web server and MySQL database, which it installs as part of the installation process. You will need to install a certificate to secure the web server; ADSelfService Plus comes with a tool to assist you with this process.

Like Specops SPPPR, ADSelfService Plus comes with the capability to send a verification code to a user's cell phone, in addition to requiring the user to correctly answer the verification questions. And like Quest Password Manager, ADSelfService Plus includes CAPTCHA in its suite of security options. However, it doesn't integrate with AD, unlike the Quest and Specops products.

ADSelfService Plus takes some time to become familiar with, partly because many of its features are three or four levels deep and partly because the language that the interface uses is easy to misinterpret. For example, you might think that the *Force user to enroll* option sets the client application to intercept the logon process and force the user to enroll. But this option actually means that the user must go through the enrollment process before

they are allowed to use the employee directory system that is built into ADSelfService Plus. So this product is a little confusing for the administrator when first using it.

### ADSelfService Plus

**PROS:** Includes a comprehensive and customizable password self-reset portal for end users, Help desk technicians, and administrators; inexpensive

**CONS:** Doesn't integrate with AD, so doesn't enforce the configured password settings unless users are in the ADSelfService Plus interface

**RATING:** ◆◆◆◆◆

**PRICE:** $995 for 1,000 users, including annual maintenance and support

**RECOMMENDATION:** If you're looking for a password management system that allows for multiple password policies but you don't have the budget for an AD-integrated product, then ADSelfService Plus is worth considering.

**CONTACT:** ManageEngine • 888-720-9500 • www.manageengine.com

## Quest Password Manager

The installation process for Quest Password Manager (which Figure 3 shows) consists of a wizard that walks you through the processes of creating a new password reset and password policy, which you then assign to a container in your AD domain as well as your AD security groups. After the wizard walks you through the password policy, password reset policy, security options, and container assignment, you'll have a good understanding of the product. In this way, the setup wizard functions as a guided tour for the administrator. Be aware that to ensure that Quest Password Manager integrates fully with your domain, you will need to install the Quest Password Manager .msi file on all domain controllers (DCs).

Of all the products I compared, this one had the most integration options. Quest Password Manager is designed to work with Microsoft Identity Integration Server or Quest ActiveRoles Quick Connect. With the latter, user information can be synchronized across AD, AD Lightweight Directory Services (ADAM), delimited text files, Microsoft SQL Server, LDAP directory services, OLE DB, Sun ONE Directory Server, an Oracle database, Novell Directory Services (NDS), IBM Resource Access Control Facility (RACF), IBM Lotus Domino Server, and the Google Apps service.

Quest Password Manager includes a Graphical Identification and Authentication (GINA) Group Policy template, which allows you to add the configuration settings for this application to your domain Group Policy. You can customize not only the template's look and position on the screen, but also the behavior of the client application. For example, you can force the use of HTTP Secure (HTTPS), statically assign the recovery center URL, or configure proxy settings.

The big difference among the products tended to be integration with Active Directory

Quest Password Manager requires a full SQL Server installation (not just SQL Server Express Edition), with SQL Server Reporting Services (SRSS) installed as well. If you don't have a SQL Server installation available, you'll need to add the price of SQL Server to your cost analysis. On the up side, you'll have SRSS to review all the information that is available in Quest Password Manager.

Another Quest Password Manager feature is the ability to assign a temporary passcode to users who haven't gone through the enrollment process. These passcodes can be configured to expire within a set amount of time. Just be careful with this feature; anyone with access to the Help desk portal can assign a passcode, then enroll and reset any account that the Quest Password Manager service account has permission to change. If you decide to use this feature, be sure to delegate the service account correctly. Otherwise, your Help desk staff might have much more access than you intended. The passcode feature is turned off by default.

### Quest Password Manager

**PROS:** Easy to install; includes a setup wizard to get your system up and working quickly; integrates with AD, allowing for assignment of separate password policies to different OUs via Group Policy settings; many options for integration

**CONS:** No significant disadvantages

**RATING:** ◆◆◆◆◆

**PRICE:** $5,000 for 1,000 users

**RECOMMENDATION:** If your company is in the market for a fully integrated password management system with the potential to integrate with several directory services, then Quest Password Manager might be the best choice.

**CONTACT:** Quest Software • 800-306-9329 • www.quest.com

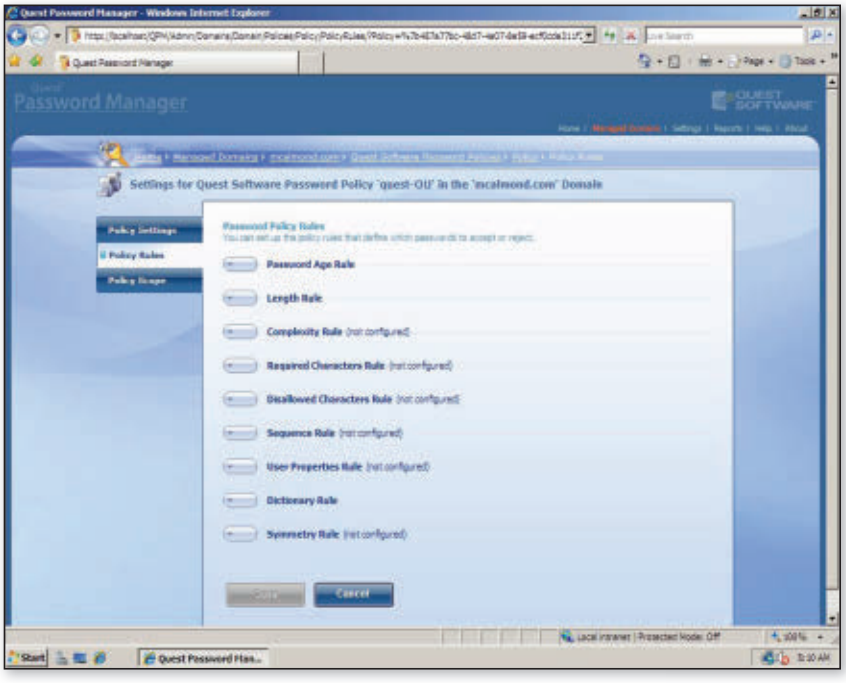


Figure 3: Quest Password Manager

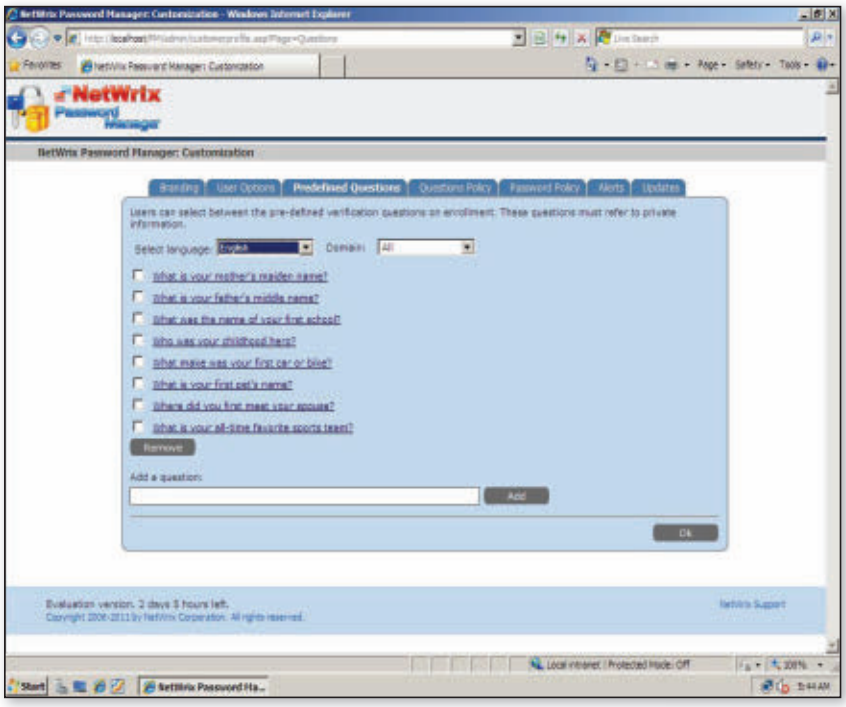


Figure 4: NetWrix Password Manager

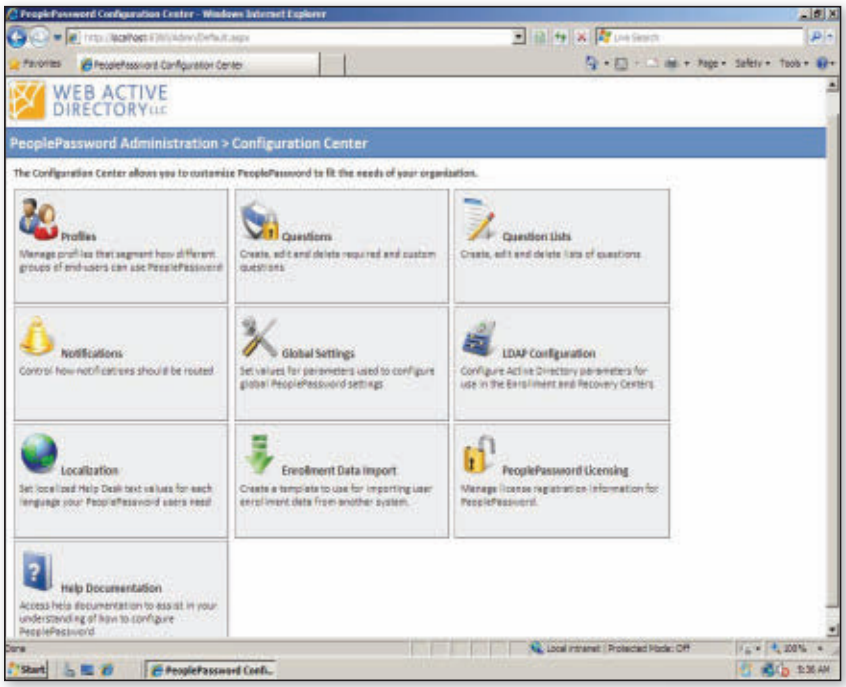


Figure 5: Web Active Directory's PeoplePassword

## NetWrix Password Manager

NetWrix Password Manager (which Figure 4 shows) installs on Windows XP Service Pack 3 (SP3) or later. After the installation, I only needed to adjust the authentication in IIS to enable Windows Authentication. After the installation, NetWrix Password Manager is very simple and intuitive. You can choose from a list of verification questions, make up your own, or allow end users to make up their own questions. Custom questions and answers can be required to be a minimum length, and all answers can be required to be unique. During the installation, the product creates an AD group called NetWrix Account Help Desk. Adding users to this group gives them the ability to use the Help desk web portal to assist other users with resetting passwords or unlocking accounts.

Unique to NetWrix Password Manager is a disconnected-mode password reset. The disconnected-mode reset enables the GINA extension on the Windows logon screen to reset a user's cached password, even when the user isn't connected to the domain. This could be a key feature for companies with large numbers of mobile users but does require the GINA extension to be installed locally.

NetWrix Password Manager also comes with a user-data import process that can be used to prepopulate the information that is needed during user enrollment, making the enrollment process easier for end users.

NetWrix also offers a freeware version for as many as 50 enrolled users.

### NetWrix Password Manager

**PROS:** Easy to use; has almost no learning curve; includes password reset capability for Google Apps

**CONS:** Does not tell users what the password requirements are during the reset process

**RATING:** ◆◆◆◆◇

**PRICE:** $6.50 per user for 150 users; significant discounts for larger numbers of users; subscription licensing is 33 percent of perpetual licensing after volume discounts

**RECOMMENDATION:** Companies with many mobile users might see the disconnected capability of NetWrix Password Manager as a major gain over the other products. Also, smaller companies might find that the 50-user freeware version is all they need.

**CONTACT:** NetWrix • 888-638-9749 • www.netwrix.com

## Web Active Directory's PeoplePassword

Web Active Directory's PeoplePassword (which Figure 5 shows) is another product that works with IIS and SQL Server. However, in this case you can use SQL Server 2005 or later, including SQL Server Express Edition if you don't already have and don't want to pay for SQL Server.

Web Active Directory has done a nice job with the enrollment process in PeoplePassword. This product comes with the ability to import all necessary user information so that you can enroll users into the system without any involvement on their part.

In addition to the core functionality that all the products provide, PeoplePassword has the unique ability to collect an alternate email address during the enrollment process. This address can then be used to send a password-reset code during the password-reset process. The ability to collect an alternate email address and send the password-reset code can be turned on or off, simply by checking a box in the password-reset profile settings. However, you can't require the verification questions to be answered before sending the password-reset email—an improvement that some companies might want before using this feature.

### PeoplePassword

**PROS:** When used with Web Active Directory's add-on PeopleEnroll product, the automatic enrollment process allows organizations to import data into PeoplePassword and complete the user enrollment processes with zero user involvement

**CONS:** Doesn't unlock and reset a password at the same time, requiring users to go through the reset process multiple times if they forget their passwords and become locked out of their accounts

**RATING:** ◆◆◆◆◇

**PRICE:** $4.50 per user per year for perpetual licensing; $0.25 per user per month for subscription licensing

**RECOMMENDATION:** If the enrollment process concerns you or you don't want enrollment to be a manual process for end users, then PeoplePassword might be your best option.

**CONTACT:** Web Active Directory • 800-747-3565 • www.webactivedirectory.com

## Product Differences

The big difference among the products I've included in this comparative review tended to be integration with AD. Two of the evaluated products—Specops Password Policy and Quest Password Manager—integrated with AD in such a way that I could assign different password policies to different OUs within a domain, even if the domain's operational mode didn't natively enable this option. In both products, an application needed to be installed on each to allow the product to intercept the password-change requests and ensure that they complied with the specified requirements before being passed on to AD.

These products enforced my password policies both when using the product interface and when using the standard change-password routine that's built into all Windows versions, from any computer in the domain, with or without a client installation.

If you're interested in more granular contrasts, see Table 1 for a comparison of all the products' core features. (I give each product one point per provided feature; for Group Policy integration, I give the product two points.)

InstantDoc ID 141113

**Nate McAlmond**
(mcalmond@gmail.com) is an MCSE, a Security+, and Network+. He's the director of IT for a growing social services agency, and he specializes in thin-client infrastructure and electronic health record implementation.

Table 1: Product Feature Set Comparison

| | Specops Password Policy and Password Reset | ManageEngine ADSelfService Plus | Quest Password Manager | NetWrix Password Manager | Web Active Directory PeoplePassword |
|---|---|---|---|---|---|
| **Usability** | | | | | |
| Web interface | 1 | 1 | 1 | 1 | 1 |
| Unlock during password reset | 1 | 1 | 1 | 1 | 0 |
| Account unlock | 1 | 1 | 1 | 1 | 1 |
| Multiple password policies | 1 | 0 | 1 | 0 | 1 |
| Multiple password reset policies | 1 | 1 | 1 | 1 | 1 |
| Multidomain management | 1 | 1 | 1 | 1 | 1 |
| Password self reset | 1 | 1 | 1 | 1 | 1 |
| Password reset via alternate email | 0 | 0 | 0 | 0 | 1 |
| Help desk password reset assist | 1 | 1 | 1 | 1 | 1 |
| Dynamic display of password requirements during password reset | 1 | 0 | 0 | 0 | 0 |
| Password requirements shown during reset | 1 | 1 | 1 | 0 | 1 |
| Multilingual support | 1 | 1 | 1 | 1 | 1 |
| Windows logon screen integration | 1 | 1 | 1 | 1 | 1 |
| Group Policy integration | 2 | 0 | 2 | 0 | 0 |
| Report tools | 1 | 1 | 1 | 1 | 1 |
| Preload enrollment data | 0 | 1 | 0 | 1 | 1 |
| Disconnected password reset | 0 | 0 | 0 | 1 | 0 |
| Configuration wizard | 0 | 0 | 1 | 0 | 0 |
| | 15 | 12 | 15 | 12 | 13 |
| **Security** | | | | | |
| Password policy strength assessment tool | 1 | 0 | 0 | 0 | 0 |
| Updating of disallowed-word dictionary | 1 | 0 | 1 | 0 | 0 |
| Cell-phone verification and notification | 1 | 1 | 0 | 0 | 0 |
| Verification question lockout | 1 | 1 | 1 | 1 | 1 |
| CAPTCHA | 0 | 1 | 1 | 0 | 0 |
| Notification of password reset sent to users or administrators | 1 | 1 | 1 | 1 | 1 |
| Password-protected enrollment | 1 | 1 | 1 | 1 | 1 |
| | 6 | 5 | 5 | 3 | 3 |
| | | | | | |
| **Total** | 21 | 17 | 20 | 15 | 16 |

# Trends in Mobile Device Security

## 7 tips for defending the smartphones and tablets in your environment

by Jeff James

It's no secret that mobile devices—from smartphones to Apple iPads to Android tablets—are appearing in the enterprise in record numbers. Some reports already indicate that more smartphones than PCs were shipped in 2011, and that trend will undoubtedly continue. The PC isn't going away, but it's being joined by a large assortment of new mobile devices in a variety of form factors.

Although all these new and powerful devices are boosting worker productivity to even higher levels, they're also introducing some thorny security problems for IT managers and security professionals. Your VP of human resources might be more effective on the road with her smartphone, but what if she leaves her unlocked BlackBerry—along with a Microsoft Excel spreadsheet listing executive salaries—in the lunch room at her HR conference? Or what about the engineer who has the detailed specs for your new product on his iPad, which he inadvertently left on the subway on the way to work? Then there are the programming interns who regularly download apps for their Android devices outside the Android market—including apps that are infected with malware and viruses.

These scenarios are all security issues, and I haven't even touched on the compliance and auditing demands that are placed on businesses and organizations that must abide by such regulations. "All of these factors are putting more pressure than ever on IT professionals, who are being pressured into allowing the use of social media tools like Facebook, who are dealing with the consumerization of IT, and who now have additional mobile devices to secure and keep track of," says Don DeBolt, director of threat research for Total Defense, which was spun off from CA earlier this year and which serves as an independent business focusing on mobile security.

### The State of Mobile Security

Judging by the mobile-security headlines of 2011, malware authors are being attracted to mobile devices in record numbers—and to the Android platform in particular. Android has emerged as the dominant smartphone OS, and with that distinction comes the attention of malware authors. There have been plenty of news reports about Android malware, ranging from infected apps in the official Android Market ("Up to 120,000 users download infected apps from Android Market," June 2011, InstantDoc ID 136942) to key-logger applications masquerading as legitimate apps ("Bogus Netflix Android App Attempts to Steal User Information," October 2011, InstantDoc ID 140886).

"Android has been a victim of its own success partly by becoming the most popular smartphone OS," says Kevin Mahaffey, CTO of mobile-security software provider Lookout. "That one argument [is] why Android is afflicted with more malware than other mobile OSs. It's also a much more popular OS in countries like China and Russia, where most malware seems to be written." Mahaffey also suggests that the ubiquity of the Java programming language makes it widely available to programmers who might consider creating malware to attack the Android OS.

Both Mahaffey and Eric Sites, chief scientist for GFI Software, draw parallels between the dominant market shares of Windows and Android as significant reasons for malware authors to target those platforms. Today's cybercriminals are just as concerned about return on investment (ROI) as any business manager would be. Why shouldn't they direct their efforts toward the mobile OS that has the most users and, logically, the best possibility of a return on their malware-coding investment?

"A lot of the trends we're following for mobile malware mirror that of the PC market," Sites says. "Ten years ago hackers were doing things for fame or the thrill of it, but now there are organized networks of

criminals out there who are attempting to control devices for more nefarious reasons, like obtaining credit card numbers, stealing corporate information, and gaining access to other sensitive data."

Sites points out that cyberattacks from groups that are funded by nation-states are on the rise. He uses an aerospace-component manufacturer as a hypothetical example: If a nation that's hostile to the United States wants to find out the specifics of a component that is used on a B-2 bomber, it can make a targeted attack—involving phishing, malware, and vulnerability exploits—to try to access that information. (This type of maintained attack is sometimes referred to as an advanced persistent threat [APT].) Sites contends that mobile devices open up even more avenues for attacker exploits, ranging from obtaining misplaced devices and using malware to redirecting email and text messages or recording and forwarding spoken conversations.

## The Social Engineering Threat

Despite spending billions on endpoint security—firewalls, antivirus software, blacklisting and whitelisting solutions, and so on—cybercriminals are still able to gain access to the most sensitive information. The culprit is social engineering, which criminals use to fool people into thinking they're replying to an email message or clicking a link from a trusted source.

Social engineering is too broad a topic to go into here (see "Protecting Yourself Against Social Engineering," January 2005, InstantDoc ID 47956, for an excellent treatise on the subject), but many experts believe that social engineering tactics are being used with greater frequency than ever before. The highly publicized attack on RSA ("RSA Reveals Details of Phishing Attack," April 2011, InstantDoc ID 136753) was caused directly by an RSA employee clicking on a file attachment in an email message that the employee believed was from a legitimate source.

## Steps to Secure Mobile Devices

Considering the current growth rates in the adoption of smartphones and tablets in businesses of all sizes, any IT pros that are tasked with managing mobile-device security have their work cut out for them. To help you get the most out of your efforts, here is a list of tips and techniques that can help you improve security for all your mobile devices.

***1. Embrace clear and direct security policies.*** If your company or organization has specific security, auditing, or compliance needs that need to be taken into account when managing mobile devices, be sure to carefully (and concisely) document the corporate security policy. It's amazing how many companies don't have a clearly communicated security policy when it comes to mobile devices, so this is a good place to start.

***2. Make employee training a must.*** With a majority of mobile cyberattacks leveraging social engineering tactics, regularly training your employees to know how to spot fraudulent apps and email messages is vital. Eric Sites also castigates app designers for contributing to the current security problems. "There needs to be more usability testing [for Android apps]," Sites says. "Many of these apps aren't the easiest for the average user to figure out, and many users may click randomly on things in order to make something happen."

***3. Quickly find or wipe lost or stolen devices.*** Several brands and models of mobile phones include support for finding lost phones or performing remote wipes. Improved tools for managing mobile devices are also available, including Odyssey Software's Athena. Athena works with Microsoft System Center Configuration Manager (SCCM) and Research In Motion's (RIM's) BlackBerry Enterprise Server (BES) to help you manage mobile devices for those platforms.

***4. Enforce good password policy.*** Good password policy is just as necessary for mobile devices as it is for desktops and laptops, so having a common-sense (and consistently enforced) password policy that applies to mobile devices is a must. A good policy enforces minimum limits for password length, character-string complexity, password expiration, and more. (For more password tips, see my blog post "Password Security Tips," December 2010, InstantDoc ID 136736.)

***5. Research apps before downloading.*** Smartphone users should be trained to carefully read about the apps that they intend to install and the permissions that those apps will require. Mobile apps that ask for access to an unusually large number of smartphone features should be instantly suspect; extraordinary requests for features can be a telltale sign that the app isn't what it appears to be. "Users should really spend some more time reading about their apps before downloading them, rather than mashing their finger down on the first thing they see," says Sites.

***6. Leverage mobile-security solutions.*** A number of security vendors—including McAfee, Symantec, Lookout, ESET, and Total Defense—have started providing security software for mobile devices. All these vendors provide differing sets of features and functionality, so your best bet is to do some research and find the solution that works best for you and the assortment of devices and OSs that you're deploying and supporting.

***7. Beware vendor FUD marketing.*** With so much ink being spilled about the security risk of mobile devices—particularly regarding Android OS security issues—it's also clear that many security vendors have jumped on the bandwagon, using fear of malware or viruses on mobile devices to drive product downloads and sales. "My reaction to this situation is 'sigh,'" says Mahaffey. "You can't fault these companies for using an extremely successful business model, but we took a different approach. We wanted to make people happy and confident about using their mobile devices without all the fear tactics." Other security software vendors that provide phone location and file backup services are also duplicating functionality that already exists in some smartphones, such as the Apple iPhone's Find My iPhone, remote wipe, and regular iCloud backup features.

InstantDoc ID 141211

**Jeff James**
(jjames@windowsitpro.com) is industry news analyst for *Windows IT Pro*. He was previously editor in chief of Microsoft *TechNet* magazine, was an editorial director at the LEGO Company, and has more than 15 years of experience as a technology writer and journalist.

# WINDOWS IT PRO VIP is

**Educational**—with FREE eLearning courses and eBooks available 24×7

**Deep**—housing over 41,000 articles on DVD and online, some exclusively for VIP members

**Broad**—solutions, tips, and tricks for any Windows or SQL Server issue that can stump you


In fact, Windows IT Pro VIP delivers more than **$1,000 of resources and expertise for just $199 a year.**

| HOW WINDOWS IT PRO VIP BEATS A SEARCH ENGINE | | |
|---|---|---|
| | **Windows IT Pro VIP Delivers:** | **Search Engines Deliver:** |
| **Reliability** | Road-tested advice from experts who put their reputation on the line | Well-meaning but potentially harmful tips in the latest Wikipedia entry |
| **Speed** | The answers you need in seconds searching by keyword, topic, or publication | Lost time spent perusing sites that have mastered search engine rankings but not the art of Active Directory or patch management |
| **Impartiality** | Authors and experts who challenge the Microsoft party line and influence industry change | Conventional wisdom touted by industry insiders afraid to tell it like it is |

**Order Online Now at windowsitpro.com/go/vip**

# SharePoint Archiving Solutions

## Archive SharePoint content to help keep your SharePoint environment running smoothly

by Caroline Marwitz

Microsoft SharePoint has become what email was a decade ago: a dumping ground for content. Companies are realizing that this content needs to be managed, secured, and—in many cases—archived.

The first two needs are obvious, but why would you want to archive SharePoint content? For three simple yet compelling reasons: data reduction (which can affect performance), governance, and compliance.

"Archiving tools . . . help you maintain the size of your content databases as well as allow for real-time version-history archiving," said Errin O'Connor, who has more than a decade's worth of experience with SharePoint and is founder and CEO of EPC Group .net. Helping achieve the goals of a governance plan is yet another reason for archiving SharePoint content. "Archiving old sites that are no longer used—this is key, as it's important to either delete or archive content that's no longer relevant," O'Connor said.

Easily retrieving that content is important as well. "You may have a project team site that was used for a project and that project is over, but in a year or two a similar project may pop up again and the project manager or team members may want to go back and restore that archived site to follow the best practices or lessons learned from that previous project," O'Connor said.

Archiving is a basic best practice in records management, but there's an even more compelling reason for some organizations. "Archiving is about compliance," said Ron Charity, a SharePoint product manager who has worked with SharePoint since 2001 and focuses on governance, information architecture, technical architecture, and operations. Compliance with industry or governmental regulations is essential for many, if not most, organizations, especially in the United States, which is home to the largest percentage of the world's lawsuits. Compliance and auditing capabilities go hand in hand with archiving. As O'Connor explained, "You can restore an archive to a site or SharePoint instance and make that data available to auditors and e-discovery activities without affecting the live SharePoint farm."

But SharePoint 2010 has the ability to manage records in place, so why would you need a third-party archiving solution? For one thing, Charity said, many organizations need a compliant archival engine (e.g., compliant with US Department of Defense—DoD—requirements). Another reason, he said, is that "enterprise records management systems scale much better due to N-Tier architecture and use of the file system for items and SQL Server for logic." Additionally, you can't beat the convenience of certain third-party products' features. "When archived data is disposed of, client systems issue certificates for legal purposes," Charity said.

What should you look for in a SharePoint archiving solution? Seamless integration with SharePoint is obvious, and vendors accomplish this goal in different ways. For example, many solutions stub the item in SharePoint and move it to the archive, whereas some solutions integrate with SharePoint at the event-handler layer to capture items. Can end users search for and access archived content in SharePoint? They'd better be able to, unless you like training them on new solutions and procedures.

E-discovery capability is useful; as part of that, so is the ability to archive all content types and data in SharePoint. Also consider how the vendor packages a solution, whether as a suite or a standalone product (only your organization's needs should determine which option is best for you). Then there are things that you won't know until you try a tool: how flexible it is, how easy it is to use, and how responsive the customer service is.

The buyer's guide table shows a sampling of SharePoint archiving vendors and the particulars of their solutions. If you're still not sold on the need for archiving SharePoint content, read the AIIM blog "The Case for SharePoint Archiving" (www.aiim.org/community/blogs/expert/the-case-for-sharepoint-archiving-ite28099s-never-too-soon-to-deal-with-old-information). Another useful blog post on SharePoint archiving and what to look for in a SharePoint archiving solution is Geoff Evelyn's "SharePoint Archiving—Defining a way Forward" (www.sharepointgeoff.com/sharepoint-archiving-%E2%80%93-defining-a-way-forward).

InstantDoc ID 141264

**Caroline Marwitz**
(sharepointeditor@penton.com) is the editor and web content manager for *SharePoint Pro* magazine, at sharepointpromag.com.

## SHAREPOINT ARCHIVING SOLUTIONS

| Company | Product | Price | SharePoint Versions Supported | E-Discovery Support | Global De-Duplication | Search | Backup and Restore | Full-Text Indexing/ Global Single-Instancing | Auditing | Policy-Based Records-Lifecycle-Management Support |
|---|---|---|---|---|---|---|---|---|---|---|
| **AvePoint** 201-793-1111 800-661-6588 www.avepoint.com | DocAve Software Platform | Contact vendor | SharePoint Foundation 2010; SharePoint Server 2010; Windows SharePoint Services (WSS) 3.0; SharePoint Server 2007 | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| **CommVault** 732-870-4000 888-746-3849 www.commvault.com | CommVault Simpana 9 | Varies depending on configuration | SharePoint Foundation 2010; SharePoint Server 2010; SharePoint Server 2007 | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| **LiveOffice** 800-374-2032 www.liveoffice.com | LiveOffice SharePoint Archive | $19.95 per account/ month (unlimited storage and retention) | SharePoint Server 2010; SharePoint Server 2007 | Yes | Yes | Yes | Yes | Yes | Yes | No |
| **Message-Solution** 408-383-0100 www.messagesolution.com | Message-Solution Enterprise SharePoint Archive | Contact vendor (government and non-profit discounts available) | SharePoint Foundation 2010; SharePoint Server 2007 | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| **Metalogix Software** 877-450-8667 www.metalogix.com | Metalogix StoragePoint | $14,995/ SharePoint front-end web server for Enterprise Edition; $7,495/ SharePoint front-end web server for Standard Edition with File Share Librarian | SharePoint Foundation 2010; SharePoint Server 2010; WSS 3.0; SharePoint Server 2007 | No | No | No | No | No | No | Yes |
| **Symantec** 650-527-8000 800-721-3934 www.symantec.com | Symantec Enterprise Vault for Microsoft Office SharePoint Server 2007 and SharePoint Server 2010 | $38/user, or about $6,800/TB | SharePoint Foundation 2010; SharePoint Server 2010; WSS 3.0; SharePoint Server 2007 | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

***Editor's Note:*** Some vendors you might expect to see in this Buyer's Guide said they didn't have a product that exactly matched the criteria or didn't respond to our requests for information about their products.

***Note:*** An expanded version of this table is available online at www.windowsitpro.com, InstantDoc ID 141264.

| | Automatic Classification and Tagging of Items | Continuous or Periodic Capture of Data | SQL Server Requirement | Location of Content Archives | Capture Methods | Option to Migrate SharePoint Data | Tiered Storage Management | Compliance-Requirement-Specific Templates | Retrieval by End User |
|---|---|---|---|---|---|---|---|---|---|
| | Yes | Both | No | Any local or network drive, file system, or cloud or SAN environment | Leverages SharePoint native forms of capture for in-place capture and routing with ability to add metadata; also integrates with third-party solutions for physical capture | Yes; can offload data to local or network drive, file system, cloud, or SAN environment | Yes | Yes | Yes |
| | Yes | Periodic | No | External BLOB Storage (EBS) store | SharePoint API | Yes; can migrate data to public or private clouds, tape storage, disk, or other media, as well as to another content-management system with Representational State Transfer (REST) interface support | Yes | Yes | Yes |
| | No | Both | No | LiveOffice cloud archive service (via Transport Layer Security—TLS—encryption); encrypted at rest with Advanced Encryption Standard (AES) 256-bit encryption | Leverages trigger-based archival methods (based on document check-in) or uses a predefined schedule | Yes; SharePoint files can be exported from vendor cloud archive to a defined destination on your network | No | No | Yes |
| | Yes | Both | No | On-premises archiving server, public or private cloud service | SharePoint API | Yes | Yes | Yes | Yes |
| | No | Both | Yes | SAN, NAS, content addressed storage (CAS), or cloud-based storage platforms; generic file systems; EMC or Hitachi solutions; or Windows Azure, Amazon Simple Storage Service (S3), or Rackspace cloud files | Remote BLOB Storage (RBS) and EBS | Yes; can externalize SharePoint content BLOBs to specific on-premises and cloud-based storage platforms through use of add-on plug-ins | Yes | No | No |
| | Yes | Both | Yes | Flat file | Crawls SharePoint and archive based on metadata, size, content type, or file type | Yes; can migrate to tiered storage, including tape and cloud options | Yes | No | Yes |

INSIGHTS FROM THE INDUSTRY

# The Usefulness (or not) of Cloud Office Application Suite Reviews

After reading the review "Google Apps vs Office 365 vs Zoho, your office in the cloud" (http://tinyurl.com/7ybf2ee), which purports to compare and contrast the competing cloud-based office application suites from Microsoft, Google, and Zoho, I wonder about the usefulness of such a report at this point in the evolution of cloud technology.

Although Google Apps has been around for several years, its competitors are still relatively new. As such, I don't think it's possible to provide a realistic assessment of just how well a broad spectrum of cloud-based products performs in production. We can certainly review Office 365 since its debut last June and discuss how well Microsoft has delivered cloud-based versions of Microsoft Exchange, SharePoint, and Lync. We might quibble at Microsoft's plans and point to the two recent public outages the service suffered from. We can compare the approach taken by Zoho to Google's or Microsoft's approach and debate whether a free version of the software is worth having. But we still can't discuss important topics that CIOs ponder as they try to decide which cloud-based suite to select from the undoubted hype that surrounds this topic.

For example, any application can have teething problems soon after it goes into production. I think this is what happened with Office 365 in its two recent outages. More importantly, CIOs are concerned with how well the application performs over a sustained period. In other words, will the oft-cited service level agreement (SLA) of 99.9% be met over a year, two years, or three to five years? And how will the support team function over that period in terms of its ability to promptly communicate with tenants if outages occur, how fast problems will be fixed, and whether the fixes will cause any effects for customers, such as requiring the deployment of a hotfix on client PCs? I think that hard, time-proven experience of an application operating under the stress of production is far more important than the initial impressions gained by a reviewer who might arrive at his or her conclusions after a week or so.

Another issue that comes to mind is how cloud-based application suites will evolve over time. The nature of technology is that it never stays static for long and its creators have an almost unstoppable urge to add new features or tweak existing code. I suspect that cloud-based applications will be no different than their on-premises cousins in the race to add new features. So, the question that swings into my mind is how well vendors will manage evolution for their tenants. For example, if we look at the current version of Office 365, we can see that Microsoft Outlook 2010 is the client that exposes the maximum feature set from Exchange Online. Features such as MailTips and retention policies remain invisible if you use Outlook 2007. These features won't make a lot of difference to many tenants, who will happily continue to use Outlook 2007. But what happens when later versions of Office 365 appear? Will Microsoft remove support for Outlook 2007 or require Outlook 2012 (or whatever it's called then) to connect to the cloud? I don't think so, but the thought of any forced client upgrade makes CIOs wince.

You can argue that Google or any other browser-based application provider has the advantage here because they don't have to deal with the complexities of a feature-rich ("fat") client. This is true because the task of the application provider is much simpler when the provider exerts full control over both client and server. But Google hasn't been immune to criticism as it has changed the UI of its applications over time, sometimes almost seemingly on a whim. Individual users will accept the change, especially if they are (like me) a consumer of Google's free services and don't have a vote. Those who pay for the service might have a different view. Companies like consistency and predictably because it makes Help desk support easier. Changes that make sense to a developer can cause confusion all around if not flagged well in advance.

I exclude task-oriented, "getting started" type articles from my criticism, such as Zac Wiggy's "Office 365 vs. Google Docs: End-User Perspective" (October 2011, InstantDoc ID 140011), because these articles actually include some useful information that tells people what they must do to even start using cloud applications. Of course, I might not agree with the author's conclusions, but that's an argument for another day.

It would be nice if we could review the cloud-based application suites in depth, but I think this might be a fool's errand. The suites haven't been around long enough for all their strengths and weaknesses to be understood and exploited. For now, we must be satisfied with the reviews that do get published, in the form of a checkmark-based feature comparison across the suites. Not ideal, but it's all we can do until we really understand what we're dealing with.

—Tony Redmond

InstantDoc ID 140707

# 10 Reasons the iPhone 4S Is Selling Like Crazy

The iPhone 4S sold more than 4 million units in the first three days after its release, and it's on track to become the best-selling iPhone ever. Why has the iPhone 4S been so successful? Here's my analysis of the factors that helped propel the iPhone 4S to the top of the smartphone sales heap.

## 1. Best iPhone Ever

Despite the mountains of hype that the Internet rumor mill generated over the imaginary iPhone 5, the iPhone 4S is still the best iPhone ever made, with specs that exceed, equal, or are very close to the best Android smartphones currently available. The iPhone 4 was beginning to look long in the tooth, but the iPhone 4S helps keep Apple competitive with other leading smartphone handsets. That said, Apple now has a host of competitors that can produce hardware that meets or exceeds what the iPhone can offer.

## 2. Upgrades that Count

The iPhone 4S might not have had many things that the Internet rumor mill predicted it would have, but it did have significant upgrades where it counted: a faster, dual-core A5 processor; up to 64GB of storage; a vastly improved camera that is arguably the best smartphone camera now available; and Siri, a voice-recognition/personal assistant app that—although still effectively in beta—is a big step forward for voice-recognition software.

## 3. Contract Math

Millions of iPhone 3GS users signed up for two-year contracts in Q4 2009, and all those contracts have expired. I definitely fall into that camp, having purchased my 3GS back in October 2009. I couldn't care less if it's called an iPhone 4S or an iPhone 5, because the 4S is a huge upgrade over my existing 3GS.

## 4. The Implosion of BlackBerry and WebOS

BlackBerry-maker RIM has been beset with falling market share, management failures, and highly publicized service outages. The technically impressive WebOS had the rug pulled out from under it after only months under the HP umbrella. The smartphone world is now increasingly looking like a two-horse race between Android and Apple, with RIM and Microsoft battling for a distant third.

## 5. Android Woes

Speaking of Android, the world's most popular smartphone OS is quickly becoming a victim of its own success. Android hardware and OS fragmentation gives both Android developers and users headaches and could potentially pose problems for the popular OS in the months and years to come. Android is also becoming the platform of choice for malware aimed at mobile devices, with more malware being developed for Android than all the other mobile OSs combined.

## 6. Simplicity

Apple excels at producing products that are easy for non-technically inclined people to use. There are dozens of stories about how the very young (and the very old) have quickly adapted to using an iPhone or iPad. Compare that with the often bewildering complexity of Android: dozens of devices from different manufacturers running slightly different version of the Android OS. Windows Phone promises to offer a happy middle ground between Apple's walled garden and Android's sprawling, chaotic jumble of an ecosystem, but it still lags behind both Android and the iPhone on hardware and software feature parity.

## 7. The iPhone Ecosystem

Apple has many iPod and iPhone customers, and they've collectively purchased millions (if not billions) of songs, ebooks, games, videos, apps, and other content from Apple. I definitely fall into this camp, with dozens of purchased iPhone apps and hundreds of songs. Granted, I could move my song and video collection to other mobile devices, but I'd have to repurchase all my apps for Android or Windows Phone 7. Who wants to go through all that hassle, especially when the iPhone 4S offers a significant phone upgrade while keeping all of your legacy investments intact? It's a strategy that has worked extremely well for Microsoft for the Windows ecosystem, and Apple is clearly taking a page from Microsoft's playbook here.

## 8. New Carriers

The iPhone 4S is the first iPhone to ever be available on Sprint, which now joins AT&T and Verizon as the three largest iPhone carriers in North America. The iPhone 4S is also available on CSpire (formerly Cellular South), the largest privately held wireless carrier in the United States. Millions of Sprint and CSpire customers now have access to the iPhone 4S, and that previously untapped customer base undoubtedly helped push the iPhone 4S into sales record territory.

## 9. The Legacy of Steve Jobs

The death of Steve Jobs was tragic news, but it also had the unintended effect of monopolizing news coverage for the weeks leading up to the launch of the iPhone 4S. Thousands of stories sang the praises of Jobs and the iPhone, which were soon followed by a rising flood of excerpts, leaks, and commentary from Walter Isaacson's excellent Steve Jobs biography. All of this Apple-focused news undoubtedly helped the iPhone 4S receive more press coverage than it would have normally.

## 10. Apple *Is* the Mass Market

What many Microsoft and Apple zealots fail to understand is that Apple is no longer just a niche computer manufacturer known for a distinctive design aesthetic. Apple is the world's most valuable tech company and a producer of millions of goods that are aimed squarely at the mass market. Microsoft fanboys overlook that Apple isn't a niche provider solely of expensive computers for self-absorbed people who only care about proving their superiority over others, and Apple zealots often don't understand that the overwhelming majority of Apple customers don't wear pajamas festooned with the Apple logo, or troll forums looking for Microsoft products to bash. The mass market wants consumer products that look good, work as advertised, and are simple to learn and use—and Apple understands and profits from that need better than any other technology company in business today.

—Jeff James

InstantDoc ID 141198

# Using OWA Over Outlook: An Experiment

When Microsoft Exchange Server 2010 launched about two years ago, one of the major themes was Outlook Web App (OWA) improvements: conversation view, MailTips, integrated presence, and IM, all available with full-featured OWA versions on multiple browser platforms. My question at the time was if OWA was now good enough to replace Outlook on the desktop (see "OWA vs. Outlook in 2010," www.windowsitpro.com, InstantDoc ID 103261).

For many, the OWA experience truly offers all you need in an email and calendaring client. But gaps still remain between OWA and the complete Outlook desktop version you get with the Microsoft Office suite. To start, let's look at some of the positives.

For creating, sending, and replying to basic email messages, OWA provides all the tools you need. Likewise, for adding meetings or reminders to your calendar, there's really nothing lacking. And it's all just about the same as doing it in desktop Outlook, including easy linking to your Contacts or Global Address List (GAL), complete with resource scheduling. You can get pop-up notifications for new mail and calendar reminders—or not, if you choose. In fact, you have a great deal of control as an end user on all your settings, including the ability to change your password, through the OWA interface.

I was able to take advantage of the integrated IM through OWA, and it works well. The Search function, which I use a lot in Outlook, works just about as well in OWA; OWA has some filtering capability, but it doesn't offer the level of fine-tuning you can get through Outlook itself.

There are a lot of good things to focus on with OWA. But let's take a look at what's missing. The first thing that struck me was that I couldn't insert an image in my email signature. After doing a little research in some Microsoft forums, I found some workarounds—but apparently these solutions aren't supported by Microsoft. And it's not that OWA can't handle images or images in a signature. For example, if I create an email in Outlook that includes an image in my signature and save it in the Drafts folder, then open the draft in OWA, the image appears just as expected.

Some companies want to use a corporate logo as part of their email signatures, typically in the form of an image file. If you're considering using OWA instead of Outlook, using the Drafts folder workaround obviously won't help you. Your options then are to go with an unsupported solution, use plain text signatures, or wait for Microsoft to add this capability to OWA.

The other major fail for me is the spell-checker. You don't get the autocorrect behavior while you type that most of us have come to rely on when using Outlook. No automatic capitalization of sentences, no changes of *ahve* to *have,* and no fixes of the other mistakes we probably make without even realizing it. You do get a red underline of (perceived) mistakes so that you can correct them yourself. You can set the spell-checker to run automatically before sending every message, and you can run it manually.

A bigger problem with this version of spell check is that you can't add words to its dictionary. Proper names (e.g., in your signature) and unfamiliar technology terms trip up the spell-checker every time. I get tired of telling it that my name is spelled correctly every time I send a message—particularly when it's in the same automatic signature that's applied to every message.

Here's a list of other problems with using OWA instead Outlook:

- Although you can sort the view with oldest messages on top or newest messages on top, OWA doesn't remember your choice from one session to the next: It always opens with newest on top.
- Flagged messages don't show up as clearly as they do in Outlook, where the whole line turns red when an item is due—which can be a big deal if you flag a lot of email messages.
- Advanced graphics and formatting options aren't available to OWA (e.g., tables, SmartArt, themes).
- No Quick Parts, which is an easy way to send email messages based on a template or to include a section of text that you use frequently and repeatedly.

Having used Outlook with all its great features is the real spoiler. I've gotten used to features that help me be more productive. But for workers who've never used Outlook, OWA will mostly likely seem like it has everything they need. Providing an excellent OWA experience is certainly in Microsoft's best interests. Although OWA with Exchange 2010 provides a great experience, power users will still find that there's room for improvement.

—B.K. Winstead

InstantDoc ID 141067

# AD INDEX

For detailed information about products in this issue of *Windows IT Pro*, visit the websites listed below.



## VENDOR DIRECTORY

The following vendors or their products are mentioned in this issue of *Windows IT Pro* on the pages listed below.



## DIRECTORY OF SERVICES | WINDOWS IT PRO NETWORK

Search our network of sites dedicated to hands-on technical information for IT professionals.
**www.windowsitpro.com**

### Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
**www.windowsitpro.com/go/forums**

### News

Check out the current news and information about Microsoft Windows technologies.
**www.windowsitpro.com/go/news**

### EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

*DevProConnections UPDATE*
*Exchange & Outlook UPDATE*
*Security UPDATE*
*SharePoint Pro UPDATE*
*SQL Server Pro UPDATE*
*Windows IT Pro UPDATE*
*WinInfo Daily UPDATE*

**www.windowsitpro.com/email**

### RELATED PRODUCTS

Custom Reprint Services

Order reprints of *Windows IT Pro* articles:
penton@wrightsmedia.com

### Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events, plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.
**www.windowsitpro.com/go/vipsub**

### *SQL SERVER PRO*

Explore the hottest new features of SQL Server, and discover practical tips and tools.
**www.sqlmag.com**

### ASSOCIATED WEBSITES

DevProConnections

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
**www.devproconnections.com**

SharePoint Pro

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.
**www.sharepointpromag.com**

### NEW WAYS TO REACH *WINDOWS IT PRO* EDITORS:

LinkedIn: To check out the *Windows IT Pro* group on LinkedIn, sign in on the LinkedIn homepage (www.linkedin.com), select the Search Groups option from the pull-down menu, and use "Windows IT Pro" as your search term.

Facebook: We've created a page on Facebook for *Windows IT Pro*, which you can access at: http://tinyurl.com/d5bquf. Visit our Facebook page to read the latest reader comments, see links to our latest web content, browse our classic cover gallery, and participate in our Facebook discussion board.

Twitter: Visit the *Windows IT Pro* Twitter page at www.twitter.com/windowsitpro.

Windows IT Pro

# Ctrl+Alt+Del

by Jason Bovberg

"Send your funny screenshots, oddball product news, and hilarious end-user stories to rumors@windowsitpro.com. If we use your submission, you'll receive a *Windows IT Pro* Rubik's Cube."

# IT Pro Cool

## PRODUCT OF THE MONTH

Expressly made for the savvy IT pro, these geeky business card cases have a tech appeal that will make you the envy of your peers. According to ThinkGeek, "Real circuit boards come in many colors and have varying degrees of wiring in them. Consequently, so do these business card holders. Revel in their uniqueness!" The case holds 10 to 15 cards (depending on the thickness of your business card). Learn more at ThinkGeek (www.thinkgeek.com/homeoffice/5d15).

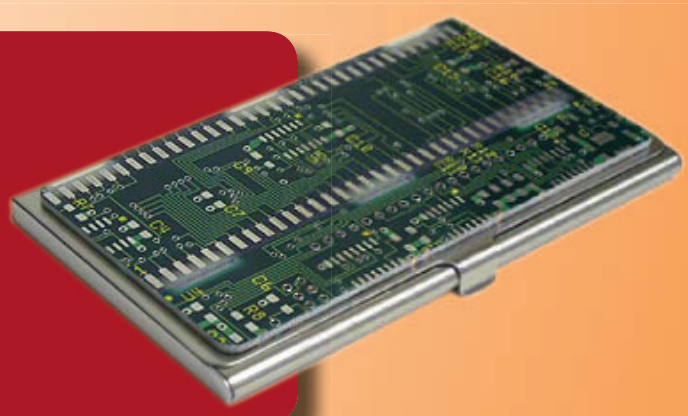

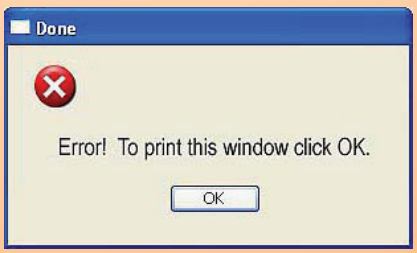


Figure 1: I guess I need a printout …

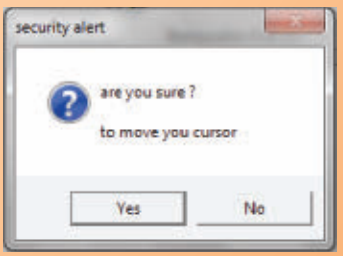


Figure 2: Uncommon security threat


Figure 3: Not really, no …

## USER MOMENT OF THE MONTH

There was one time I was working on a user's computer, and while I waited for an operation to complete, I was idly spinning the cursor around the screen. Perhaps you do the same thing. Curious, the user asked, "What are you doing there?" I said, seriously, "Sometimes you have to spin the mouse clockwise to wind it up. It's not something you need to do often, but I'll do it when I'm working on something else, just to take care of it." I expected the user to laugh and roll his eyes, but he swallowed it with a "Huh!" I shared a laugh with the rest of the team. A few days later, one of my guys was working on the same machine. The user caught him moving the cursor around, too. He asked, "Wait, I thought you were supposed to spin the cursor clockwise?" Without missing a beat, the tech replied, "Sometimes, they get wound up too tight, and you have to unwind them."

—Chris

January 2012 issue no. 209, *Windows IT Pro* (ISSN 1552-3136) is published monthly.  *Windows IT Pro,* 748 Whalers Way, Fort Collins, CO 80525, (800) 793-5697 or (970) 663-4700. Sales and Marketing Offices: 748 Whalers Way, Fort Collins, CO 80525. Advertising rates furnished upon request. Periodicals Class postage paid at Fort Collins, Colorado, and additional mailing offices. **POSTMASTER: Send address changes to *Windows IT Pro,* 748 Whalers Way, Fort Collins, CO 80525.** SUBSCRIBERS: Send all inquiries, payments, and address changes to *Windows IT Pro,* Circulation Department, 748 Whalers Way, Fort Collins, CO 80525. Printed in the USA.

MARCH 26-29, 2012 • MGM GRAND • LAS VEGAS, NEVADA

CLOUD CONNECTIONS | WINDOWS CONNECTIONS | Microsoft® Exchange CONNECTIONS

SharePoint CONNECTIONS | SQL Server CONNECTIONS

# TAKE THE JOURNEY INTO 2012

## SEE NEW PRODUCTS & TECHNOLOGIES FROM MICROSOFT


SOLUTION ARCHITECTS, SYSTEM AND MESSAGING ADMIN, INFRASTRUCTURE ARCHITECTS, NETWORK ADMIN, CTOS, IT MANAGERS AND DIRECTORS

***KEEP YOUR COMPETITIVE EDGE!***

Connections is excited to be hosting one of the SQL SERVER 2012 LAUNCH EVENTS

## 2012 LAUNCH
CONFERENCE AND EXPO

powered by Microsoft® & SQL Server

### KEYNOTES

**MARK MINASI**
*MINASI RESEARCH AND DEVELOPMENT*

**QUENTIN CLARK**
*MICROSOFT*
CORPORATE VICE PRESIDENT, DATABASE SYSTEMS GROUP, MICROSOFT SQL SERVER

**SCOTT GUTHRIE**
*MICROSOFT*
CORPORATE VICE PRESIDENT, SERVER & TOOLS BUSINESS

**STEVE FOX**
*MICROSOFT*
DIRECTOR, MICROSOFT CONSULTING SERVICES

### CONFERENCE ADVISORY BOARD

**KIMBERLY L. TRIPP**
*SQLSKILLS.COM*

**PAUL S. RANDAL**
*SQLSKILLS.COM*

**MICHAEL OTEY**
*WINDOWS IT PRO & SQL SERVER MAGAZINE*

**CHRIS AVIS**
*MICROSOFT*

**SEAN DEUBY**
*PENTON MEDIA*

**DON JONES**
*CONCENTRATED TECHNOLOGY, LLC*

**LEE MACKEY**
*DELL CONSULTING*

**JIM MCBEE**
*ITVICOS SOLUTIONS*

**PAUL THURROTT**
*WINDOWS IT PRO MAGAZINE*

**HAROLD WONG**
*MICROSOFT*

*...and many more!*

**FIND US!** facebook.com/winconnections

**FOLLOW US!** twitter.com/winconnect

**Register NOW** Take advantage of early bird registration and hotel discounts.

REGISTER TODAY! www.WinConnections.com • 800.438.6720 • 203.400.6121

Microsoft

Files accessible from virtually anywhere.
Access from multiple devices.
A shared desktop in the cloud.
**It all works together.**

Introducing Microsoft Office 365. Collaborate in the cloud with Office, Exchange, SharePoint, and Lync videoconferencing. **Starting as low as $10 per user per month. Begin your free trial now at Microsoft.com/office365**

Scan tag with a smartphone to learn about the Office 365 free trial.
Download the free scanner app at http://gettag.mobi

Microsoft® Office 365

January 2012

## The Essential Guide to

# Least Privilege: A Use Case-Based Approach

**You have many different ways to better manage least privilege in your Windows infrastructure; be sure to select the solution most targeted for your specific needs**

**By Jan De Clercq**

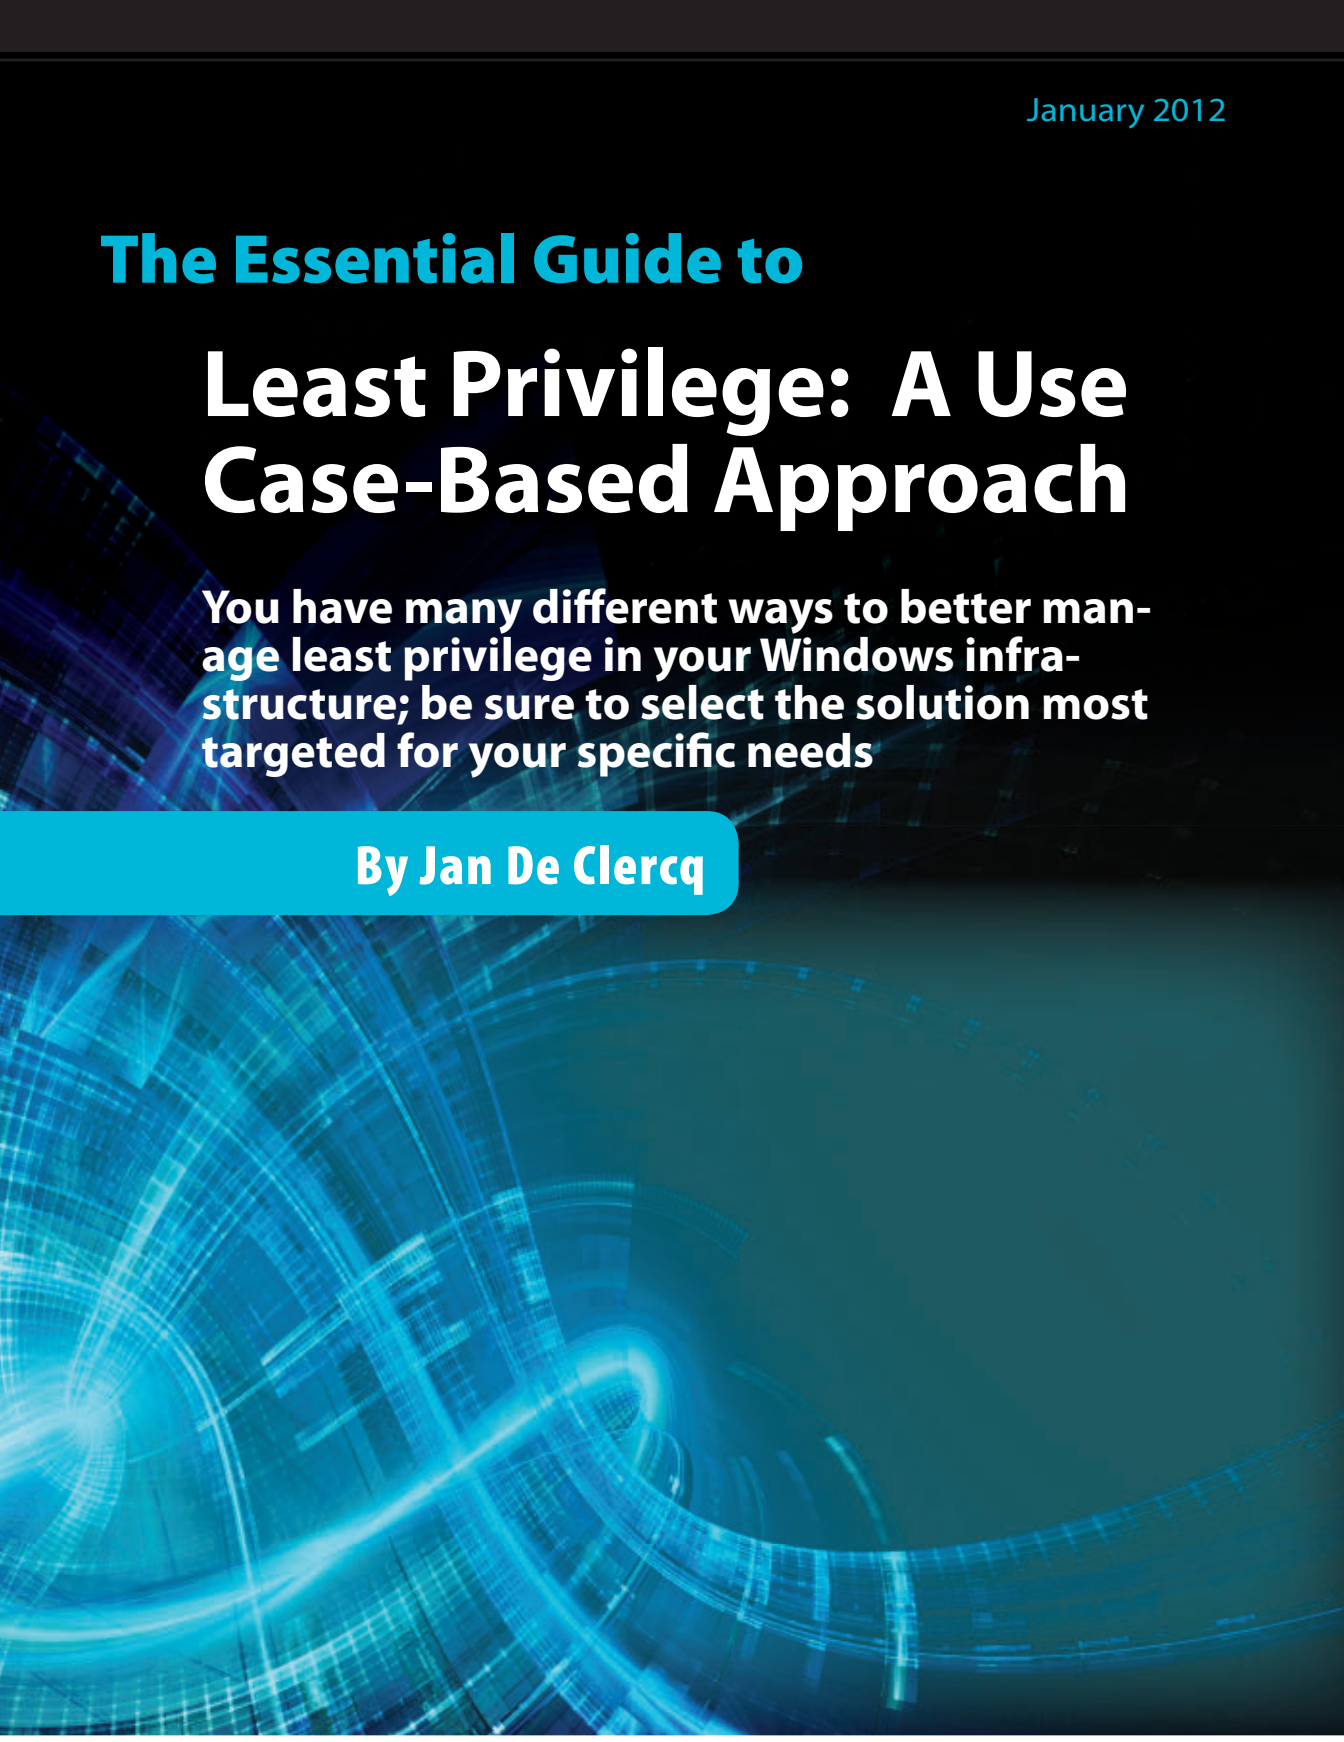

VIEWFINITY

*SPECIAL ADVERTISING SUPPLEMENT TO WINDOWS IT PRO*

An administrator accidentally downloads malicious code while surfing the Web. A Windows developer writes code that requires administrator privileges to work properly. These dangerous practices violate one of security's most fundamental principles: the principle of least privilege. This principle states that you should give a user or a piece of code only the privileges it needs to do the job—nothing less, and certainly nothing more. Malicious code can do much more harm when it can executes in the security context of a highly-privileged account, and highly privileged processes can do much more harm when they are compromised or simply buggy.

Least privilege has long been a well-respected and supported principle in the UNIX world, but Microsoft started taking it seriously only with the release of Windows 2000 and Windows XP. Support for the least privilege principle was a key security theme of Microsoft Vista when Microsoft introduced User Account Control (UAC). Starting with Windows Vista Microsoft fundamentally redesigned the way Windows supports least privilege. In Windows 7 Microsoft provides further optimizations to UAC.

This Essential Guide explains which tools administrators and users can use in Windows today to honor least privilege. It also explains how third-party applications like Viewfinity's Privilege Management (PM) solution can complement the least-privilege controls embedded in Windows to further lock down access to the Windows platform and its resources.

## Least Privilege in Windows XP

Windows XP provides two tools to help you honor least privilege; unfortunately, few Windows users and administrators know of their existence. You can either use Fast User Switching (FUS) or the RunAs tool to run Windows processes and applications in the security context of a non-administrator account. The RunAs (runas.exe) command line utility is the easiest way to switch to the security context of another user when you're using a domain-joined machine. On standalone machines you would use Fast User Switching (FUS).

FUS is an XP Professional Edition and XP Home Edition feature that allows for multiple, simultaneous interactive logon sessions on a Windows computer. FUS lets users easily switch between logon sessions without logging off or closing running applications. FUS availability is subject to several limitations: it is available only on machines that are not joined to a domain, have the Welcome logon screen enabled, and don't have a Graphical Identification and Authentication (GINA) Module replacement (for example a custom GINA that enables biometric authentication or another strong authentication method) installed. Furthermore, the user mustn't have enabled Offline Files support.

FUS is not enabled by default: you can enable it from the Control Panel User Accounts applet, using the "Change the way users log on or off" option. To effectively use FUS you must then create an account that has computer administrator privileges, and next, create a user account that has limited privileges. You can then use the user account for day-to-day work such as surfing the Web, email, or Instant Messaging (IM). Switch to the security context of the administrator account when you need to install programs or run programs that require administrator privileges. When you finish such jobs, log off the administrator account and switch back to the user account. Because this account has reduced privileges, you don't necessarily need to log off the user account each time you leave your computer. Simply locking the screen is usually sufficient.

The RunAs utility is rooted on the Secondary Logon service, which lets you start logon sessions that use other credentials within the current logon session. The Secondary Logon service is installed by default in Windows 2000, XP, and later systems and starts automatically when the system boots. When you are logged with a user account and want to start the Microsoft Management Console (MMC) in an administrator security context, you can type the following RunAs command at the command line:

```
runas /user:administrator mmc
```

RunAs then prompts you to enter the administrator's password, as Figure 1 shows. If the password is correct, the utility will start an instance of the MMC.

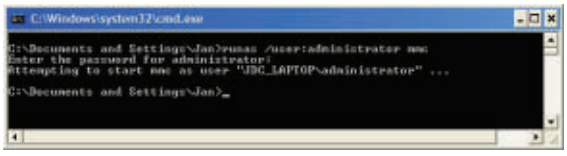


**FIGURE 1: Using RunAs on Windows XP**

You can also run RunAs from Windows Explorer. To do so, you must right-click an executable's or shortcut's icon and select "Run as" to open the Run As dialog box. This dialog box presents options to select alternate credentials. For icons that don't offer the Run as option in the regular context menu (for example Control Panel applets), hold down the Shift key while right-clicking the icon. You can also make RunAs the default action for a particular executable or shortcut by opening that executable's or shortcut's Properties, clicking Advanced, and selecting the "Run with different credentials" check box.

## XP Least Privilege Limitations

The problem with XP and earlier Windows OSs is that they require a lot of user discipline to honor least privilege, from both users and administrators. For example, XP administrators who want to honor least privilege must create two accounts: a simple user account and an account with administrator-level privileges. Then, they must have the discipline to use the user account to perform their day-to-day work—surfing the Web, reading email, working with Microsoft Office documents—and switch to their administrator-level account only to perform administrative tasks.

It is true that in XP Microsoft made it easier to switch between logon sessions (FUS and RunAs are accessible from the Windows UI), but it's still far easier for administrators to use one privileged account to do all their day-to-day work. Using a single account also means the administrative user must remember and maintain only one set of credentials. The bigger problem is that not only XP administrators but also simple XP users typically use a single account with administrator-level privileges. Working with a limited account in XP can be a frustrating experience because the OS doesn't allow limited user accounts to perform simple administrative tasks such as changing a system's time zone settings, installing additional fonts, or changing power-management

options. And for these reasons, administrators often grant simple users admin-level privileges.

The bottom line is that average XP users and administrators who value ease of use more than security and who don't want to switch back and forth between two user accounts leave their system open to a wide range of attacks. And let's be honest: How often have you used the built-in local administrator account to log on to your standalone XP workstation?

## Introducing User Account Control

Starting with Windows Vista, every account—including the built-in administrator account and other administrator-level privileged accounts—initially has only limited user privileges. During logon sessions, users can elevate their privileges to the administrator level when necessary, as I will explain later. This functionality is what Microsoft refers to as the standard User Account Control (UAC) behavior.

Administrator accounts with initially limited user privileges is possible thanks to a change that Microsoft has made in Vista—specifically, the process of creating access tokens for privileged-account users. An access token contains a user's privileges and is attached to a user-logon session. When a privileged-account user logs on to Vista, the OS creates two tokens: a filtered token and a full token. The filtered token contains only the user's limited-account privileges and is the user's default token during the logon session. The full token contains all the user's privileged-account privileges. Vista attaches the full token to the filtered token when the user needs to perform an administrative task or launch an application that requires privileged access.

Another fundamental least-privilege-related Vista change is that Microsoft has redefined what a limited account can and can't do. For example, Vista lets a limited-account user change the system's time and time zone settings, change display properties, install additional fonts, and change power-management options. This modification essentially removes the need for the Power Users group, which Microsoft has eliminated from Vista. Examples of Vista tasks that still require a privileged account are software installation and disk repartitioning.

Microsoft also clearly shows what actions require administrator-level privileges and which don't in Vista. All operations that require administrator-level privileges are marked with a shield icon, as Figure 2 shows. Figure 2 shows Vista's Date and Time Properties dialog box. Note that only the *Change Date and Time* button requires administrator-level privileges; any user can change the time zone settings. Administrative buttons that are marked with a shield icon are also called unlock buttons. In typical enterprise environments, only Help desk personnel will use unlock buttons—for example to, when they need to control desktops remotely. In typical home environments, only parents will use the unlock buttons—for example, to make configuration changes for children. With Vista, a Help desk operator or parent can, for example, unlock a Control Panel applet without an employee or child needing to log off first. Normal users can't use an unlock button because they don't know the password of privileged accounts.

When a user clicks an unlock button, selects an action that requires administrator privileges, or starts an installation program that requires administrator privileges, Vista will behave differently depending on whether the user is a limited-account or privileged-account user. If the user is a limited-account user, Vista prompts for alternate administrator credentials; if the user has administrator-level privileges, Vista prompts for consent. Figures 3 and 4 show the dialog boxes that Vista displays when a limited-account user and a privileged-account user, respectively, attempt to change the system date and time.

In the dialog box that Figure 3 shows, a limited-account user can select one of the system's administrator-level accounts and enter the administrator credentials, or can cancel the privilege-escalation action. The dialog box allows only for password credentials; it

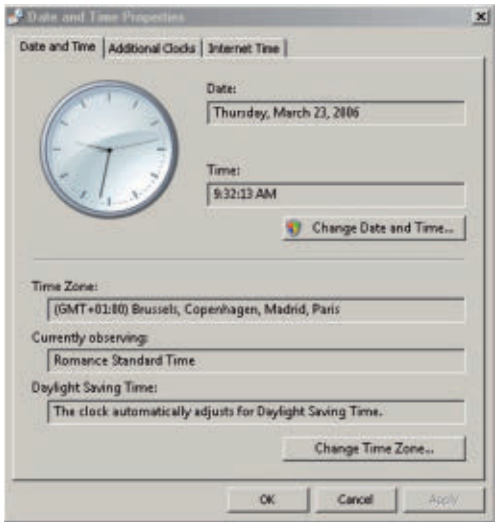


**FIGURE 2: Windows Vista Date and Time Properties dialog showing actions requiring administrator-level privileges**

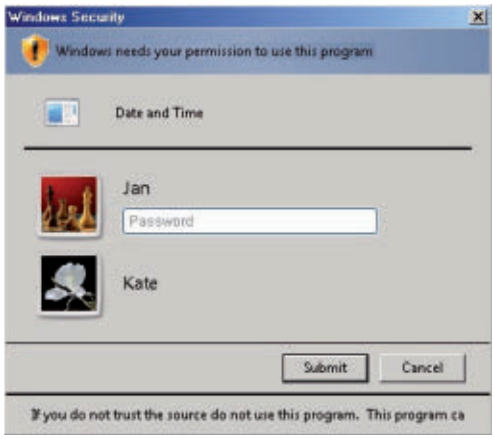


**FIGURE 3: Windows Vista UAC dialog for limited-account users**

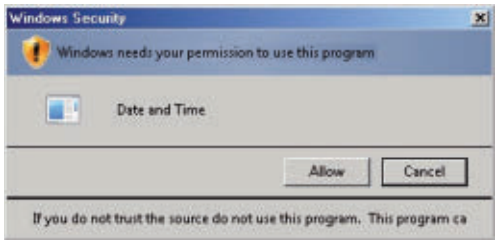


**FIGURE 4: Windows Vista UAC dialog for privileged-account users**

doesn't provide a mechanism to, for example, retrieve smart card–based administrator credentials.

In the dialog box that Figure 4 shows, a user with administrator-level privileges can express his or her consent by simply allowing or cancelling the action. Vista's Group Policy Objects (GPOs) include settings that let administrators change the default Vista prompt behavior for limited-account and privileged-account users. You can find the UAC-related GPO configuration settings—which begin with the words "User Account Control"—in the Security Settings\Local Policies\Security Options GPO container.

Microsoft refers to this privilege-based behaviour as Administrator Approval Mode (AAM). Thanks to AAM, users and administrators can honor least privilege in a single logon session. Switching back and forth between limited-account user and privileged-account user logon sessions is no longer necessary.

Credential or consent dialog boxes, such as the ones that Figures 3 and 4 show, can appear multiple times during a user's logon session, appearing whenever a user chooses an action marked with a shield icon. Vista doesn't remember previous elevations to administrator-level privileges. The elevated privileges that are linked to a particular task (for example the installation of a new software package) automatically expire when the task is finished. For this reason, AAM significantly reduces the Vista attack surface. AAM also presents an important advantage for administrative users, who can now—rest assured—perform their day-to-day work as regular users and switch to administrative privileges only as necessary or when prompted.

## Applications and UAC

Not only do certain user and administrator actions require administrative privileges, also certain applications may require higher privileges to function properly.

Similarly to administrative actions, applications that require administrator-level privileges are marked with the shield icon, on top of their standard application icon in the Vista interface. Independently of whether an application is marked as requiring runtime administrator-level privileges, users can request to start an application in the security context of an administrator account: by selecting "Run as administrator" in the program's context menu. (This menu appears when you right-click the program icon.) This option is also available in the context menu of shortcut icons.

Users can also configure an application to always run with administrative privileges: Simply select the "Run this program as an administrator" check box on the Compatibility tab of the application's Properties. Both limited-account and privileged-account users can access these options. Note that selecting "Run this program as an administrator" doesn't change Vista's AAM behaviour: The system will still prompt limited-account users for administrative credentials and privileged-account users for consent.

Vista and later Windows OSs support several mechanisms to actually "mark" an application as requiring runtime administrative privileges:

- Based on a given application's properties, Vista automatically classifies certain applications as requiring runtime administrative privileges. For example, setup or installation applications are automatically marked as requiring administrative privileges.
- During application development, a developer can mark an application as requiring runtime administrative privileges. He or she does so in the application's manifest file.
- An administrator can install an application-compatibility shim on the machine that marks an application as requiring runtime administrative privileges. With this approach, an administrator can let a legacy application start with administrative privileges without making code changes.

To test your legacy applications for privilege concerns, Microsoft provides special checks that are embedded in the Application Verifier tool.

Vista's UAC can also deal with applications that aren't marked as requiring runtime administrator privileges and that must write to a registry or file system location that requires administrator-level access privileges. This functionality is possible thanks to the UAC file-system and registry virtualization features. UAC will transparently create the data an application requires in a file system or registry location that's accessible by using limited-account user privileges. UAC virtualization also automatically redirects applications to the virtual locations when applications must retrieve or write data.

## Windows 7 UAC Changes

In Windows 7 Microsoft increased the number of tasks that a limited-account user can perform and that do not prompt for administrator approval. For example, limited-account users can now install updates from Windows Update and reset network adapters without receiving a UAC dialog. In Windows 7 Microsoft also disables the built-in Administrator account by default. It can also not log on to a computer in Safe Mode.

A very visible Windows 7 UAC change is the new User Account Control Settings dialog that is available from the User Accounts Control Panel applet. It allows a privileged-account user to configure the UAC experience ranging from "Always notify" to "Never Notify". Windows 7 includes 4 different UAC configuration levels. Windows Vista offers only two options: UAC is either on or off.

Finally, in Windows 7 Microsoft also provide additional local security policies to enable a local administrator to change the behavior of the UAC messages for privileged-account users and for limited-account users.

## Windows Least Privilege Limitations

UAC is a great step forward for better honouring least privilege in Windows Vista and later Operating Systems (OSs). Unfortunately Microsoft will never support UAC on Windows XP. On this platform you are bound to user discipline for properly using FUS or RunAs for honouring least privilege.

An important UAC limitation is that the default tasks that a limited-account user can do and that Microsoft exempts from UAC are fixed and cannot be customized. The default exemptions simply may not be sufficient for all organizational needs. For example, organizations may want to give certain users the ability to install certain ActiveX controls in their browser. Windows does not provide a granular management mechanism for defining what exact tasks should be exempted from UAC.

Windows UAC also does not support an exception mechanism for supporting the automatic elevation of privileges to administrator level for certain applications. Many legacy applications simply cannot live without having administrator privileges. Also for some applications the workaround of file system and registry virtualization simply does not work.

Besides UAC and the other mechanisms outlined above Microsoft also provides other solutions for limiting what users can do on their Windows machines. You can use Group Policy Object (GPO) settings for locking down the user desktop. But GPOs can only be applied to domain-joined machines and are not a solution for locking down standalone machines or machines that rarely connect inside the corporate network. Today many users connect to corporate networks with personal Windows devices that are never joined to a domain and thus always exempted

from GPO policies—a trend commonly referred to as the consumerization of IT. And we should point out that even domain-joined machines may bypass recent GPO settings it they do not log on to AD for a longer period of time.

## The Viewfinity Solution

Viewfinity provides an advanced and flexible privilege management and application control solution that deals with the issues outlined in the previous section.

To best meet an organization's requirements, Viewfinity provides three flexible server-side implementation options for deploying its privilege management and application control solution. The Viewfinity Privilege Management can be implemented through the Viewfinity SaaS/Cloud platform, via an on-premise Viewfinity server in your private cloud, or as an extension to Group Policy, enabling policies to be managed through the standard Windows Group Policy management tools.

On the client side Viewfinity Privilege Management (PM) uses a special agent that runs on the user desktop and that is available for Windows XP (SP3), Windows Vista, Windows 7, Windows Server 2003 and Windows Server 2008. The agent caches privilege policy settings it gets from a Viewfinity privilege management (PM) server.

When using the Viewfinity SaaS/Cloud platform deployment option the Viewfinity PM server is basically a web service that is hosted on the Internet, in the cloud. Whenever the desktop connects to the Internet, it pulls updated policies from and provides status information to the Viewfinity server. Thanks to the caching the agent can always enforce policies even if it is not connected to the network. In this scenario, the client agent and management server do not rely on AD Group Policy Objects (GPOs), and as a result the Viewfinity privilege policies can be enforced regardless of the desktop connection state to the AD or the network. This means that the policies can also be applied to machines that are not joined to an AD domain or forest.

Viewfinity's second deployment option uses an in-house ViewfinityPM server inside your organization. This option can be leveraged by organizations with very large user and machine populations. Both the Viewfinity SaaS and in-house server deployment options are referred to as the "server-based" model.

As a third deployment option for organizations that want do want to leverage GPOs for enforcing privilege management policies,Viewfinity provides a "server-less" option that does not require a dedicated policy management server and that enforces privilege management policies using a set of GPO extensions. Since GPOs can only be applied to domain-joined machines, you lose in this scenario the support for privilege management on non-AD clients.

Viewfinity PM allows an organization to have very granular control over what exact processes can run with elevated rights on the user desktop. You can define policies that enable a certain application that is launched from a given machine to run with elevated rights. When the applications starts, you can enforce that there will be no UAC dialog that is presented to the user and that a standard installation process is used. You also use Viewfinity PM to grant granular administrative rights for installing certain ActiveX controls based on digital signature from a specific publisher, URL, or class ID.

Viewfinity PM also includes granular policy controls for various administrative and maintenance tasks: for example, the creation of a backup, the initiation of a disk fragmentation, or simply changing appearance and personalization settings such as display, language, date and time settings. Figure 5 shows the central Viewfinity privilege management policy console.

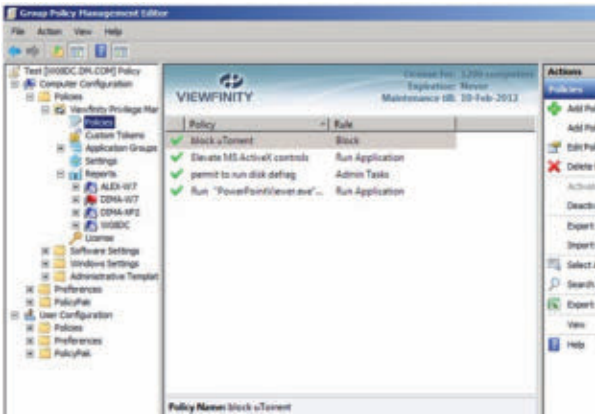


**FIGURE 5: Viewfinity privilege management policy console**

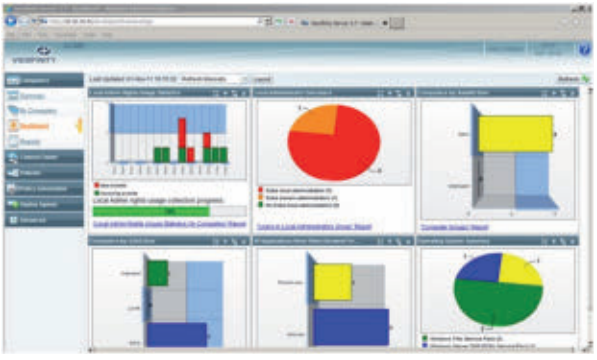

**FIGURE 6: Sample Viewfinity privilege management reports**

Bundled with the Viewfinity PM solution is also a whitelisting feature for configuring what applications can run on user desktop. The granular Viewfinity policies can be targeted to AD organization units (OUs) or to specific groups that the administrator defines from the Viewfinity console. In addition to whitelisting, Viewfinity PM can also block specific applications to offer an additional layer of control.

Viewfinity PM includes centralized reporting tools. The Viewfinity agent feeds privilege management policy compliance data to the central server that the system administrator can use for checking how and when privilege management restrictions are applied on user desktops. The central server also logs statistical information such as the most frequently blocked applications and what privilege elevation policies are used the most. It also supports alerting to inform the IT security team about privilege management actions taken were taken and that may cause servers to be less secure, such as removal of the firewall software or the disabling of an anti–virus program. Figure 6 shows some sample Viewfinity privilege management reports as they are available on the PM server.

Viewfinity PM can also integrate with configuration management tools such as Microsoft System Center Configuration Manager (SCCM). This integration provides better insights on the status of the privilege management policy on user machines and on the privilege elevation requests from users. This integration requires an add-on component that must be installed on the SCCM server. More information on the SCCM integration can be found at http://www.viewfinity.com/Products/PrivilegeManagement/SCCM.aspx.

Last but not least, Viewfinity PM offers tools to assess what exact applications require administrator-level privileges, to discover what local accounts are member of the built-in

Administrators group, and to automatically create privilege elevation policies based on this information. The Viewfinity Application Admin Rights Analysis feature can find out what exact applications, processes, and administrative actions will require administrative permission. Viewfinity offers a free Local Admin Discovery tool to discover user accounts and groups that are members of the local Administrators built-in group on the computers in your Windows domain.

## Use Cases for Least Privilege Implementation

Transitioning to a least privileged environment on user platforms is a best practice for all organizations. In earlier sections we reviewed the out-of-the box tools that Microsoft provides to restrict permissions on user platforms. But often these tools aren't flexible enough to meet a company's specific privilege management requirements.

In this last section, we will present three specific use cases that can drive the adoption of the Viewfinity Privilege Management (PM) solution. Viewfinity PM can benefit organizations that want more flexible least privilege management implementation options, easier compliance management for least privilege rules, and better control over the least privilege of mobile workers.

### Flexible Implementation Options

The Viewfinity PM solution allows for a flexible least privilege implementation that exactly fits an organization's needs. Viewfinity can be deployed using either a server-less or a server-based model.

The Viewfinity server-less model is implemented as an extension to Group Policy Objects (GPOs), enabling policies to be managed through the standard Group Policy Management tools. It also leverages AD and GPOs to deliver privilege management policies to user platforms. With this model there is no need to add additional servers or change an organization's existing infrastructure. The server-less model is a typically the best implementation model for organizations that only need to manage the privileges of local, on-premise users (there is no need to manage mobile or remote workers) and that do not have extensive least-privilege reporting requirements.

The Viewfinity server-based model leverages the Viewfinity cloud SaaS platform to manage privileges. It integrates with but is not dependent on AD and GPOs. In this model, Viewfinity provides the server environment. Because the server-based model doesn't require the installation of a local PM server, it allows organizations to implement privilege management very rapidly. To support the server-based implementation model in very large environments and for organizations that want more control over the PM server, its data and reports, Viewfinity also provides an on-premise server option.

You can find a detailed feature comparison between the server-less and the server-based Viewfinity PM deployment model at www.viewfinity.com/Support/PMFeaturesTable.aspx.

### Easier compliance management

Today most organizations must adhere to different industry and/or governmental regulations that impact IT security. Often these regulations impose the principle of least privilege. Compliance requirements make it imperative that organizations have automated compliance validation tools that include reporting and auditing features that can collect and report on the compliance and least privilege status on a large number of end points.

The Viewfinity PM solution can significantly ease compliance management through its centralized and advanced reporting capabilities and its integration with the Microsoft System Center Configuration Manager (SCCM) configuration management solution.

### Better control over mobile workers

A key challenge for today's organization is the ever-increasing need to better manage mobile and remote users. As we pointed out above there are limitations when you plan to manage privileges of mobile workers using the standard Windows AD/GPO tools. To properly deal with the privilege management of mobile and remote users, organizations need a solution such as Viewfinity PM that operates independently of AD/GPO.

Viewfinity PM can also deal with a distributed workforce that is spread out across different AD domains in different AD forests. In this case it is not possible to manage using a single AD GPO infrastructure. Viewfinity's SaaS or on-premise server-based PM can centrally control least privilege on Windows computers that are defined in different AD domains and forests using a simple HTTPs connection.

## A Complementary Solution

The built-in Windows least privilege features complemented with the Viewfinity privilege management (PM) solution can offer a more complete coverage to help you and your users to better honor the principle of least privilege. The granular and flexible privilege management policies of Viewfinity PM allow organizations to more easily reach least privilege and compliance objectives on the Windows platform of their local, remote and mobile users without sacrificing user productivity or increasing the helpdesk support call volume.

Through the use of automated settings, policies can be established and applied for multiple dimensions of configurable, logical groupings: departments, applications, end users, connectivity status, time of day, and more as depicted in Table 1. Building automated policies, combined with Viewfinity's Privilege Management features, offer a flexible yet secure approach to ensuring all your laptops and desktops automatically adhere to your corporate regulations, including compliance needs as they relate to USGCB/FDCC, PCI Compliance, and SOX—as well as other desktop-level control procedures.

**Table 1 – Use Case:** Sample Privilege Policies by Role and Responsibility

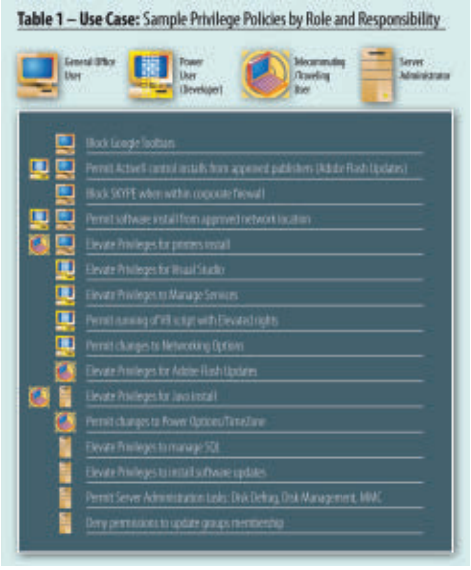


**Jan De Clercq** is a member of HP's Technology Consulting IT Assurance Portfolio team. He focuses on cloud security, identity and access management, architecture for Microsoft- rooted IT infrastructures, and the security of Microsoft products. He is author of *Windows Server 2003 Security Infrastructures* (Digital Press) and co-author of *Windows Security Fundamentals*. You can reach him at jan.declercq@hp.com

January 2012

# The Essential Guide to Optimizing Windows Server Workloads


SPECIAL ADVERTISING SUPPLEMENT TO WINDOWS IT PRO

Articles focusing on the optimization of IT operations are typically replete with buzz-words and tend to focus on abstract business-related topics that commonly feel like they're better suited to bean-counters than IT pros.

The problem, however, is that when IT departments don't implement a concrete strategy for ensuring efficiency, it's entirely too easy to get bogged down in expensive, tedious, day-to-day troubleshooting tasks and operations that can negatively impact morale, decrease end-user productivity, and stymie overall business agility and efficiency.

## Modern IT Trends and Challenges

Unless modern IT organizations take a strategic, proactive, approach to systems management, they'll end up mired in the trenches as time, energy, and effort is directed at reacting to repetitive tactical problems – which, in turn, means that IT departments can quickly become a liability instead of being the strategic technical asset that IT professionals want their organizations to be.

This is especially true in today's IT climate where IT professionals have so many competing demands and priorities for their attention, focus, and efforts. Consider, for example, just some of the larger or more prominent trends and concerns that IT departments have to address today:

- **Increased Need for Additional Security Measures and Auditing Requirements.** With more and more security breaches consistently making headlines today, it goes without saying that security is a huge concern for modern enterprises. Not only do IT professionals need to ensure the proper configuration and hardening of the data and resources under their jurisdiction, but compliance with increasingly complex regulatory requirements adds additional auditing overhead. Together, the need for additional security and auditing yields additional IT complexity along with increased demand for server compute and storage needs.
- **Increasing Emphasis on Business Intelligence.** Another common need within modern enterprises today is an ever-increasing emphasis and demand for Business Intelligence (BI) services and features. Importantly, this need appears to be snow-balling as more and more businesses adopt their own initiatives, benefit from the results – creating, effectively, an arms-race within their respective industries due to the competitive advantages afforded by increased amounts of performance data, reports, and intelligence. This, in turn, causes more demand for server computing power, additional need for storage, and results in more complexity, provisioning, and management overhead for IT professionals.
- **Increased Need for Collaboration and Interaction.** As enterprises strengthen their competitive advantage through use of Business Intelligence and other strategic initiatives, they are increasingly turning to the need for increased collaboration, interaction, and the ability to share mission-critical data. To address this need, many organizations are turning to Microsoft SharePoint – which, in turn adds additional complexity and need in the form of compute, storage, and security overhead.
- **Increased Demand for Storage.** Underlying the demand for increased collaboration, Business Intelligence, regulatory auditing, and the need to keep pace with ever larger applications and solutions is a constant demand for larger, faster, and more redundant storage.
- **Persistent Demand for Compute Power and Scalability.** In addition to the performance demands needed to keep pace with additional auditing, business intelligence, and collaboration initiatives, IT departments need to address the fact that business applications not only get more and more complex over time, but more numerous as well. Moreover, as more and more applications and complexity make their way on to the stage, IT professionals are expected to be able to more quickly and more efficiently spin-up, or provision, new servers and hosts for these increasingly complex applications and solutions. This, in turn, places increasing demand upon IT professionals who need to build, configure, and troubleshoot an increasingly larger number of bigger and more complex hosts that are doing increasingly larger amounts of processing and work.
- **Constant Pressure to Decrease IT Costs.** Yet, despite the need to juggle so many competing concerns and demands, one thing remains steady or constant within the IT landscape: the need to cut costs and reduce operating expenses. And the biggest problem with cost-cutting is that it's so difficult to quantify and a measure IT costs and expenditures. For example, while calculating depreciating capital expenditures for hardware can

be easy (for accountants), trying to separate capital costs from recurring maintenance costs can be problematic when engineers are spending large amounts of time troubleshooting new infrastructure or new initiatives. More importantly, however, is the fact that while something as intangible as 'opportunity' cost can be much harder to quantify – it's abundantly clear that IT departments can either significantly help or hinder the ability of businesses to pursue opportunities that increase business agility and overall competitiveness.

## Recent IT Trends for Addressing Challenges

Happily, several important technical trends have emerged in recent years thanks to their ability to help IT departments address the need for increased performance and throughput while simultaneously addressing the need to cut costs and decrease complexity.

### Server Virtualization and Consolidation

Chief among these emerging trends in recent years has been server virtualization – which has helped businesses combat problems associated with server sprawl that, in turn, were caused by means of attempting to address issues associated with application complexity and instability through the use of physical isolation. Server virtualization has, in turn, helped facilitate large consolidation initiatives within enterprises as a means of helping IT departments better stay ahead of increased demand for computing power while decreasing overall IT footprint and manageability costs and considerations.

### Virtual Desktop Infrastructure Initiatives

In a similar vein, a recent and obvious extension to the benefits afforded from server virtualization has been in an increased push towards achieving greater flexibility and manageability of desktop computing needs through the use of highly optimized virtual desktop infrastructures. In effect, virtual desktop infrastructures allow IT professionals to leverage strengths typically associated with server-room or back-end competence as a means for taming the desktop.

### Cloud Computing

Still in formative stages and definitely subject to a degree off buzz-word-overload, Cloud Computing does provide practical benefits as IT departments are learning about the benefits of having powerful hosting platforms at their command in order to better address workload spin-up and self-provisioning needs for highly-demanding consumers such as quality-control and developer divisions. Similarly, in terms of High Availability and systems redundancy, Cloud Computing – especially in the form of well-defined and well-secured private clouds provide significant benefits and flexibility.

## A Better, More Strategic, Approach

However, as beneficial as these emerging trends have been to the IT industry in general, they still impose additional overhead and manageability considerations in larger, enterprise, environments. Stated differently, virtualization (in any of its forms) is a tool – not a panacea that magically makes IT problems go away. In fact, while problems stemming from server sprawl are commonly addressed by means of virtualization, it's also possible for problems with server sprawl to spiral out of control when virtualization is thrown into the mix – especially when virtualization is leveraged as a mere tool or component – instead of being part of an over-arching strategy of proactively addressing manageability concerns.

### The Key to Managing Windows Server Workloads

When IT organizations allow focus to consistently drift to pain-of-the-moment concerns it becomes nearly impossible to provide applications, end-users, and solutions with the power, agility, and throughput needed to enable IT to become a strategic business asset. It also becomes nearly impossible to reign in operational costs.

Accordingly, IT departments need infrastructure, components, and manageability solutions that work in unison – towards a common, unified, goal of decreased management overhead. Otherwise, they risk becoming bogged down with tactical concerns and in troubleshooting one-off operations that don't contribute to strategic goals or the improvement of overall business agility.

To this end, there's been a significant push within enterprise IT circles within recent years towards more comprehensive management solutions—as a means of allowing administrative actions and efforts to scale better. And while a number of compelling manageability solutions are available today as a means to help IT departments transition to a more strategic approach to systems administration, one of the most powerful examples of a solution that facilitates strategic management is Cisco's Unified Computing platform – which blends storage, compute, virtualization, and unified fabric management into a single, cohesive, whole that not only enables easier management but which also ensures better overall throughput, performance, and agility.

**Figure 1:** *Cisco's Unified Computing System (UCS) unifies network, compute, storage, and virtualization into a single cohesive system.*

Engineered from the ground up as a means of addressing modern IT workloads, Cisco's Unified Computing System (or UCS), represents a next-generation approach towards a unified data-center that favors strategic vision over wasted cycles dealing with non-unified systems and components.

## Strategic IT in Action—a Case Study

As a case-study of how a consolidated approach to IT management can provide comprehensive benefits, consider the example of CareCore National, LLC –a leading Health Benefits Management Company which saw an increase from 10 million covered patients to 70 million covered patients in just 5 years. As CareCore's IT staff struggled to keep up with such rapid growth, systems administrators found that more and more of their time was consumed with troubleshooting and basic systems administration. Not only was this a problem for CareCore's IT staff, but it also negatively impacted end-users when some operations that should have executed in sub-second time-frames began to regularly take up to 10 or 15 seconds to execute. Moreover, not only were tactical problems causing pain for IT staff and end-users, but it was also causing CareCore to divert focus from business initiatives needed to integrate addition scientific advances and benefits into their business flow and core offerings.

In a bold move, CareCore decided to adopt a new strategy for systems management that would enable them to focus less time on reactive, or tactical, concerns and direct more attention to addressing core business needs—as a means of addressing end-user demands for increased throughput and performance. To drive this needed strategic change, CareCore turned to a solution provided by Cisco, VMware, and EMC in the form of their jointly-managed VBlock infrastructure initiative. This, in turn, resulted in a powerful virtualization solution built atop 168 Cisco Unified Computing System (UCS) Blade Servers handling over 400 virtual machine workloads running atop VMware vSphere – along with new Cisco networking infrastructure and a generous allotment of EMC VMax storage as a means of not only hosting typical business workloads such as Active Directory, File Servers, and Exchange, but also as a means of providing a powerful platform for CareCore's SQL Server and custom Application Server workload needs as well.

Results were telling and paint a clear picture of the benefits provided by explicitly adopting an IT management strategy as opposed to merely allowing IT resources to be caught up in a seemingly endless race to address one-off problems and issues. For example, in addition to being able to virtualize 98 percent of their server infrastructure, CareCore is now able to deploy entire, complex, production environments in a matter of literally minutes—as opposed to the 8 to 12 hours it once took. Similarly, application workloads are now highly optimized—to the point that application response times are back down to the sub-second response times desired which, in turn, has made it possible for CareCore's call-center agents to handle a whopping 20 percent more calls per day than previously possible.

More importantly, however, is the fact that a new, strategically managed server platform has enabled engineers to transition to a point where they're now able to focus 80 percent of their time on addressing new business initiatives and services – whereas they weren't even able to dedicate 50 percent of their effort to capital improvements previously. Likewise, CareCore's CTO also boasts that increased IT agility has enabled the deployment of new business initiatives that once took up to roughly 6 months to now transpire in many cases in just a matter of weeks—meaning that CareCore's IT staff has been able to dramatically help improve overall business agility in terms of responding to market needs and meeting customer demands.

Obviously, it goes without saying that an Essential Guide sponsored by Cisco's line of next-generation servers would highlight the performance benefits afforded by deployment of UCS Blade Servers. However, it's important to note that while CareCore's server platform is decidedly more robust than it was, some of the biggest benefits that CareCore has enjoyed have actually stemmed from in a dramatic increase in overall IT manageability that has correspondingly translated into better business agility along with decreased administrative costs.

Stated differently, CareCore National's success highlights how transitioning IT operations towards a more strategic and pro-active approach to management represents an essential component of streamlining IT costs, managing growth and end-user expectations and optimizing IT workloads.

## The Benefits of Consolidated Infrastructure

The benefit of a consolidated approach to IT infrastructure is that individual components are designed to be part of a greater whole—instead of just being designed for a single purpose or task. As such, converged infrastructure solutions enable a "synergistic" benefit where the benefits of the entire system working as a whole can easily outweigh the benefits afforded by each, individual, component or consideration.

Accordingly, the remainder of this Essential Guide outlines just some of the ways in which IT shops today can benefit from the increased control, manageability, performance, and cost-benefits afforded by a more strategic, and synergistic, approach to systems management by means of using smarter infrastructure and components.

## Compute Resources

In order for consolidated infrastructure to be viable, it has to be powerful – meaning that it needs the ability to address the toughest workloads and possess the capacity to scale on demand. Furthermore, above and beyond the need for raw performance and scalability, consolidated compute infrastructure needs improved manageability benefits if it's going to help IT departments cut costs.

### Raw Compute

In terms of raw compute power, Cisco's UCS line of servers has managed to set 45 world performance records—both with their high-end rack-mount C-Series servers and their B-Series blade servers.[1] Cisco UCS servers also feature support for Intel's latest line of Xeon E7 processors—which sport up to 10 cores per processor.

Likewise, because compute depends upon more than just processing power, Cisco UCS Servers also boast the ability to provide up to 512GB of RAM in a single 2-socket blade, or up to 1TB of RAM in their UCS C460 rack-mount servers—making Cisco UCS servers fully capable of handling today's toughest workloads.

### Balanced Growth and Pricing Features

While Cisco UCS Extended Memory Technology (EMT) facilitates the ability to address large amounts of physical memory, it also provides substantial cost-savings benefits for RAM-hungry workloads. For example, with a single, 2-socket, UCS B440 blade-server, EMT allows four physical 8GB DIMMS to be represented as a single 32GB logical DIMM. Not only does this provide a 60 percent cost-savings compared to competing platforms, but it also enables high-memory 2-socket configurations that don't require as much capital expense up-front or energy, footprint, and cooling costs over time as would be required by a comparable 4-socket offering from other vendors. This, in turn, means that IT shops can more easily address many of the memory-hungry workloads in their environment (such as SQL Server, Exchange, and high-density/memory-hungry virtualization hosting) without having to switch to a larger and more expensive chassis.

### Provisioning Benefits

While powerful and scalable hardware is a must for staying abreast of business demand, IT organizations also need to be able to efficiently provision, configure, and optimize compute hosts in order to truly remain agile and meet business needs. To this end, Cisco's Service Profiles represent a policy-based way for IT administrators to easily configure UCS compute resources for various workloads and responsibilities. For example, as part of an initiative using Cisco UCS blade servers as part of an integrated VDI solution that drastically increased IT responsiveness and agility along with enabling care providers to spend at least 45 more minutes per day with patients, the IT staff at Seattle Children's Hospital[2] was able to leverage Cisco's Service profiles as a way to let even inexperienced IT engineers add new blade servers when needed – as Cisco Service Profiles ensure that once a blade-server is 'plugged in' it becomes completely configured within seconds thanks to policies defined in advance.

Cisco's service profiles are also highly complimented by means of Cisco's Validated Designs (CVD) Program—a vast library of specialized reference architectures and best-practices configurations that target industry-specific best-practices configurations for common workloads such as SQL Server, Exchange, SharePoint, File Servers, App Servers, Virtualization hosts, and so on. Consequently, by leveraging these extensive resources, IT professionals are better able to 'dial in' the specific configurations needed to help them pro-actively and aggressively optimize their Windows Server workloads.

### Hardware Refresh Considerations

Since unified, top-end, infrastructure provides higher performance than non-unified hardware, it's able to achieve longer periods of service. This, in turn, translates into less-regular capital expenses for hardware and also means that IT professionals don't have to address concerns associated with provisioning, configuration, licensing, and the orchestration of hardware refreshes as regularly.

As an example of how less-frequent hardware refreshes can be beneficial to business and IT organizations, Forrester Research[3] conducted an in-depth interview with five Cisco UCS platform customers to determine cost-savings benefits. Among a host of different financial benefits afforded by adoption of the UCS platform, Forrester found that these five companies were able to save a combined total of $186,663 by means of refresh cost avoidance alone—reinforcing the way in which strategic benefits continue to deliver substantial benefits over time.

## Storage Access and Performance

Modern IT storage is largely able to keep pace with business storage demands primarily through the

use of high end Storage Array Networks, or SANS—which consolidate the capacity and throughput of large numbers of disks into a unified solution that enables increased performance and manageability while decreasing overall operating costs. However, as increasing numbers of server workloads need increasing amounts of storage—both in terms of throughput and capacity—and as greater numbers of workloads need fully-redundant (i.e., multi-path) access to storage, IT professionals are left with increasingly larger numbers of switches, cables, and interconnects to manage. Even worse, of course, is the fact that this array of switches, cabling, and interconnects needed to manage 'storage networks' closely mirrors and shadows existing networks for normal connectivity and bandwidth needs—to say nothing of the additional complexity added by other cabling needed for out-of-band management connection needs as well.

With so many distinct, yet overlapping, cabling and networking needs, it's easy to see why the definition, assignment, configuration, and maintenance of 'networking' infrastructure can consume so much time and effort for IT organizations. Accordingly, a key component of so-called "cloud" architectures of late is their heavy reliance upon network "Fabric" or highly converged interconnects that decrease the amount effort needed to configure, troubleshoot, and manage disparate communications networks by seamlessly converging them into a single set of highly flexible, highly performant, and highly manageable interconnects.

### Cisco's Unified Fabric and the Cisco UCS Server Platform

Cisco's Unified Fabric is a high-performance, virtualization-aware, network fabric that converges, or unifies, SAN and LAN traffic over a single set of cables to simultaneously ensure massive IO scalability while reducing total cost of ownership by as much as 45 percent thanks to decreased infrastructure and management costs. Unsurprisingly, one of the key benefits of Cisco's UCS line of servers is its native support for Cisco Unified Fabric—making it easier for engineers to enjoy the benefits of Unified Fabric when using UCS servers than when using other hardware platforms. With Cisco's Unified Fabric, enterprises can:

- Access the first and largest portfolio of 10 GB Fibre Channel over Ethernet (FCoE) and Fibre Channel (FC) adapters in the industry.
- Converge IO and Ethernet infrastructure to decrease overall cooling, rack, power, and floor-space requirements for solutions that demand significant IO throughput.
- Ensure massive scalability—even with large and complex workloads.
- Seamlessly integrate with Cisco Data Center Network Manager (DCNM) for simplified provisioning (of Physical and Virtual machines) as well as enhanced monitoring and up-time management.

## Conclusion

With infrastructure that works in complimentary fashion—instead of working as disparate or distinct units that focus on a single task—IT professionals can spend less time troubleshooting and more time helping businesses adopt greater overall agility and competitive advantage.

As such, while Cisco does sell powerful, high-end, enterprise-caliber servers, it's important to note that Cisco is not just aiming to sell more hardware. Instead, Cisco is aggressively tackling the problems that IT professionals regularly deal with today in terms of decreasing manageability overhead and costs through a strategic vision that focuses on creating a unified and cohesive server platform that can easily scale, adapt, and meet today's toughest business challenges.

Furthermore, as addressed by some of the examples cited in this Essential Guide, it's clear that Cisco's ability to use its status as an industry leader is helping Cisco to redefine IT operations today—to the point of being able to drive significant optimizations and improvements in the overall management of Windows and other Server workloads.

## Call to Action

To learn more about Cisco's datacenter capabilities for Microsoft's Windows Server, Hyper-V and workloads such as SQL Server, SharePoint, and Exchange please visit www.cisco.com/go/microsoft .

[1] See www.cisco.com/en/US/prod/ps10265/industry_benchmarks.html for numerous benchmarks and records including the Cisco UCS C460 M2 Rack Server benchmarks from August 2011.

[2] See www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/case_study_c36_675367.pdf for more information about how Cisco UCS in conjunction with VMware and EMC (through their joint VBlock initiative) were able to help IT staff at the Seattle Children's Hospital stop wasting 90 percent of their time chasing repetitive, one-off, work station issues that were hampering IT optimization efforts.

[3] See www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/total_economic_impact_forrester_research.pdf for more details about ways in which UCS customers saved over $1,345,000 in just three years.